Pattern Coding Meets Censoring: (almost) Adaptive Coding on Countable Alphabets

Anna Ben-Hamou* Stéphane Boucheron† Elisabeth Gassiat‡

arXiv:1608.08367v2 [cs.IT] 1 Sep 2016

Abstract

Adaptive coding faces the following problem: given a collection of source classes such that each class in the collection has non-trivial minimax redundancy rate, can we design a single code which is asymptotically minimax over each class in the collection? In particular, adaptive coding makes sense when there is no universal code on the union of classes in the collection. In this paper, we deal with classes of sources over an infinite alphabet, that are characterized by a dominating envelope. We provide asymptotic equivalents for the redundancy of envelope classes enjoying a regular variation property. We finally construct a computationally efficient online prefix code, which interleaves the encoding of the so-called pattern of the message and the encoding of the dictionary of discovered symbols. This code is shown to be adaptive, within a $\log\log n$ factor, over the collection of regularly varying envelope classes. The code is both simpler and less redundant than previously described contenders. In contrast with previous attempts, it also covers the full range of slowly varying envelope classes.

Keywords: countable alphabets; redundancy; adaptive compression; regular variation ; patterns ; minimax;

I. Introduction

Lossless coding consists of mapping in a one-to-one way finite sequences of symbols $x_{1:n} = x_1, ..., x_n$ (also called messages) from a finite or countably infinite alphabet $\mathcal{X}$ to binary sequences so as to minimize the expected length of codewords. The mapping (or code) is not only assumed to be one-to-one but also uniquely decodable (any concatenation of codewords can not be parsed into codewords in more than one way). Sequences are assumed to be generated by a stationary and memoryless source, defined as a product probability measure P^n, where $P \in \mathcal{M}_1(\mathcal{X})$ is a probability measure on the alphabet $\mathcal{X}$.

Throughout the paper, we identify uniquely decodable codes with probability distributions (thanks to the Kraft-McMillan inequality) and the notation $Q_n \in \mathcal{M}_1(\mathcal{X}^n)$ is used to denote *coding distributions*. Under this identification, the length of the codeword assigned to a sequence $x_{1:n}$ is at most $\lceil -\log Q_n(x_{1:n}) \rceil$ (here and throughout the paper, $\log$ denotes the base-2 logarithm while $\ln$ denotes the natural logarithm) [15]. The *expected redundancy* of the coding distribution Q_n (the expected number of additional bits beyond the entropy used by Q_n to encode a sequence generated by the source) then corresponds to the Kullback-Leibler divergence (or relative entropy) between P^n and Q_n: $D(P^n, Q_n) := \sum_{x_{1:n} \in \mathcal{X}^n} P^n(x_{1:n}) \log \frac{P^n(x_{1:n})}{Q_n(x_{1:n})}$. In decision theoretic language, redundancy is also called cumulative entropy risk [26].

When facing a source class $\mathcal{C}$ with a common alphabet $\mathcal{X}$ rather than a single source, the *universal coding* problem consists of finding codes that perform well over the entire source class $\mathcal{C}$. For a given class $\mathcal{C}$ of sources on $\mathcal{X}$, we define $\mathcal{C}^n = \{P^n : P \in \mathcal{C}\}$, the class of product distributions induced by $\mathcal{C}$ on $\mathcal{X}^n$. The performance of a code (a coding distribution Q_n over $\mathcal{X}^n$) with respect to a source class is quantified by the *maximal expected redundancy* defined as

$$\overline{R}(Q_n, \mathcal{C}^n) = \sup_{P \in \mathcal{C}} D(P^n, Q_n)\,.$$

The infimum of $\overline{R}(Q_n, \mathcal{C}^n)$ over all Q_n, is called the *minimax redundancy* of $\mathcal{C}^n$:

$$\overline{R}(\mathcal{C}^n) = \inf_{Q_n \in \mathcal{M}_1(\mathcal{X}^n)} \overline{R}(Q_n, \mathcal{C}^n)\,.$$

Minimax redundancy quantifies the hardness of universal coding with respect to $\mathcal{C}^n$. A source class $\mathcal{C}$ is said to have a non-trivial *minimax redundancy rate* if $\overline{R}(\mathcal{C}^n) = o(n)$. In the language of mathematical statistics, universal coding is the information theoretical counterpart of density estimation under cumulative entropy risk [See 26, for a scholarly presentation of this correspondence].

There are situations where universal coding is achievable and where minimax redundancy rates are precisely known. The simplest and most important setting consists of stationary memoryless sources over finite alphabets ($|\mathcal{X}| = k$) where

$$\overline{R}(\mathcal{C}^n) = \frac{k-1}{2} \log \frac{n}{2\pi \mathrm{e}} + O_k(1)\,,$$

as demonstrated in a series of paper that culminate with [13, 47, 48, 14]. (The notation $O_k(1)$ indicates that the upper bound may depend on k). Moreover it is known that in this setting optimal coding distributions are mixtures of product

* Université Paris Diderot, Sorbonne Paris Cité, CNRS UMR 7599 Laboratoire Probabilités et Modèles Aléatoires 75013 Paris
† Université Paris Diderot, Sorbonne Paris Cité, CNRS UMR 7599 Laboratoire Probabilités et Modèles Aléatoires 75013 Paris & DMA Ecole Normale Supérieure rue d'Ulm, 75005 Paris
‡ Laboratoire de Mathématiques d'Orsay, Univ. Paris-Sud, CNRS, Université Paris-Saclay 91405 Orsay

distributions. Dirichlet distributions with parameters $1/2, \ldots, 1/2$ over the $k-1$ dimensional simplex define the so-called Krichevsky-Trofimov coding distributions which form an important building block in our construction [See 23, for details].

This suggests that, while redundancy grows only logarithmically in the message length, it grows rapidly with the alphabet size. Recent results provide us with a refined picture when alphabet size grows with message length [42, 20, 49, 38].

If the considered collection of sources is too large, minimax redundancy may turn out to be trivial (scaling linearly with message length). In other settings, the source class may be the union of smaller classes with widely differing minimax redundancy rates (for example sources defined by finite context trees over a finite alphabet have redundancy rates that depend on the shape of the context tree).

Adaptive coding then considers an appropriate, more general approach. Assume that the excessively large collection of sources is the union of smaller subclasses and that, for each subclass, the minimax redundancy rate is non trivial and a good universal coder is available. Is it then possible to engineer a single coding method that performs well over all subclasses in the collection?

Let $(\mathcal{C}(\alpha))_{\alpha\in A}$ be a collection of source classes. A sequence $(Q_n)_{n\geq 1}$ of coding probabilities is said to be *asymptotically adaptive* with respect to $(\mathcal{C}(\alpha))_{\alpha\in A}$ if for all $\alpha \in A$

$$\overline{R}(Q_n, \mathcal{C}(\alpha)^n) = \sup_{P\in\mathcal{C}(\alpha)} D(P^n, Q_n) \leq (1+o_\alpha(1))\overline{R}(\mathcal{C}(\alpha)^n) \tag{1}$$

as n tends to infinity. If the inequality (1) holds with a factor other than $(1+o_\alpha(1))$ (that may depend on the source class $\mathcal{C}(\alpha)$) larger than 1 to the right, then we say that there is adaptivity *within* this factor. Context-Tree-Weighting [46] provides an example of an adaptive code with respect to sources with bounded or unbounded memory over finite alphabets [11, 21].

The notion of adaptivity comes from mathematical statistics [44], and adaptive coding is much related to competitive estimation [41]. It is also known as twice-universal coding [37] or hierarchical universal coding [31].

When facing countably infinite alphabets, adaptive coding is a natural problem: when the alphabet is infinite, a theorem due to Kieffer [30] entails that there does not exist a code Q_n such that for all distribution P over $\mathcal{X}$, $\frac{D(P^n,Q_n)}{n} \to 0$, that is, there is no universal code for the entire class of stationary memoryless sources over $\mathcal{X}$ and for this class the minimax redundancy rate is trivial.

This result has not deterred investigations on coding against countable alphabets. There are indeed sensible motivations for investigating such a framework: Orlitsky and Santhanam [32] observe that applications of compression techniques (text, image and so on) usually involve alphabets that are very large (words or even groups of words) and that the traditional transformation of words into letters and then into bits hides useful dependencies. Kieffer's theorem prompted at least two different approaches:

i) Orlitsky and Santhanam [32] separate the description of strings over large alphabets into two parts: description of the symbols appearing in the string, and of their pattern, the order in which the symbols appear. They redefine the performance criterion by focusing on compressing the message's *pattern* [32, 35, 22, 20, 39];
ii) investigating the redundancy on smaller source classes that satisfy Kieffer's condition. The so-called envelope classes investigated in [8] form an example of such classes [See also 1, 7, 10].

This paper pursues both lines of research: we deal with the collection of so-called envelope classes, but the adaptive code we introduce and investigate will turn out to be a pattern encoder. In contrast with [32, 35], we handle simultaneously dictionary and pattern encoding.

Definition 1 (ENVELOPE CLASSES AND ENVELOPE DISTRIBUTIONS)**.** Let f be a non-increasing mapping from $\mathbb{N}_+$ to $(0,1]$, with $1 < \sum_{j\in\mathbb{N}_+} f(j) < \infty$. The *envelope class* $\mathcal{C}(f)$ defined by the function f is the collection of distributions which are dominated by f: $\mathcal{C}(f) := \left\{P : \forall j \in \mathbb{N}_+,\ p_j \leq f(j)\ \right\}$. Define $\ell_f = \min\left\{\ell \geq 1, \sum_{j=\ell}^{+\infty} f(j) \leq 1\right\}$. The associated *envelope distribution* F is defined as $F(k) := 1 - \sum_{j>k} f(j)$ if $k+1 \geq \ell_f$, and $F(k) := 0$ otherwise. The *envelope probabilities* $(f_j)_{j\geq 1}$ are defined by $f_j = F(j) - F(j-1)$.

Envelope classes provide a framework where the search for adaptive coding strategies is feasible Falahatgar et al. [20].

The contribution of our paper is two-fold:

i) we use Poissonization arguments introduced by Acharya et al. [1] to derive tight bounds on the minimax redundancy of envelope classes. Karlin's framework (see Section II-B) then provides a tractable interpretation of those bounds, and allows us to readily obtain asymptotic equivalents when the envelope satisfies a so-called regular variation property. We thus fill the gap between previously proved upper and lower bounds for minimax redundancies of such classes;
ii) we construct a simple coding scheme, which we call the Pattern Censoring Code (PC code). This code performs an online encoding or decoding of a sequence of symbols. It pertains to the family of censoring codes described in [8, 6, 7, 10]: in the PC code, first occurrences of symbols are censored, that is they are encoded using a general purpose encoder for integers (such as Elias encoding [19]) and are implicitly inserted into a dictionary; symbols that have already been inserted into the dictionary are translated into their rank of insertion in the dictionary and then fed to a Krichevsky-Trofimov encoder that works on an alphabet that matches the current size of the dictionary. The Krichevsky-Trofimov encoder actually performs a variant of pattern coding as introduced in [32]. We show that the PC code is adaptive, within a $\log\log n$ factor, over the

collection of envelope classes whose decay is regularly varying. We thus get, with a single code, almost adaptivity over envelope classes for the whole range of possible regular variation.

The paper is organized as follows. In Section II, we lay down the notation and introduce Karlin's framework and regular variation. The main results are stated in Section III. In Section IV, we provide a sketch of the derivation of the redundancy bounds by combining the Poissonization techniques of Acharya et al. [1] and Karlin's framework. In particular, when the envelope f has a so-called *light tail* (a regular variation index equal to zero), we show that $\overline{R}(\mathcal{C}(f)^n)$ is asymptotically equivalent to

$$\mathsf{R}_f(n) := \log(\mathrm{e}) \int_1^n \frac{\vec{\nu}_f(1/t)}{2t} \mathrm{d}t \,,$$

where $\vec{\nu}_f(x) = |\{j \geq 1,\, f_j \geq x\}|$. This characterization is a powerful generalization of the tight bounds that have been established for memoryless sources over finite alphabets. The latter can be regarded as envelope classes where $\vec{\nu}_f(x) = k$ for some k and all small enough x. Indeed $\frac{k-1}{2}\log n$ scales like $\mathsf{R}_f(n)$ with respect to both k and n. The quantity $\mathsf{R}_f(n)$ also provides an equivalent of the minimax redundancy for envelope classes defined by log-concave envelopes (such that $f(k)f(k+2)/f(k+1)^2 \leq 1$ for all $k \geq 1$) that was characterized in [7]. Up to a constant factor, the minimax redundancy for envelope classes defined by heavy-tailed envelopes (with positive regular variation indexes, as investigated in [10]) also scales like $\mathsf{R}_f(n)$. This is the content of Theorem 5.

The bounds on minimax redundancy rates in Karlin's framework also suggest a universal coding strategy for each envelope class. In words, when encoding the n^{th} symbol in the message, it seems sensible to handle symbols with probability larger than $1/n$ (frequent symbols) differently from symbols with probability smaller than $1/n$ (rare symbols). The probability of frequent symbols can be faithfully estimated while the probability of rare symbols can barely be estimated from the message. This is widely acknowledged in the blossoming literature concerning estimation of discrete probability distributions, see [45]. The censoring code approach described in [8, 6, 7, 10] explores that kind of path. Here, we combine pattern coding and censoring so as to manufacture a simple encoder that achieves adaptivity within a $\log\log n$ factor with respect to all regularly varying envelope classes (except the extreme case of *very heavy-tailed* envelopes, corresponding to an index equal to 1, see Section II-C). The Pattern Censoring Code is described in Section V, and Section VI provides a modular analysis of its redundancy. Proofs are gathered in Section VII.

II. Background and notation

From now on, the message alphabet $\mathcal{X}$ is the set of positive integers $\mathbb{N}_+ := \mathbb{N} \setminus \{0\}$, and the source probabilities will be denoted by their probability mass functions $(p_j)_{j\geq 1}$.

A. Occupancy counts

Define N_n^j as the empirical count of symbol j, that is the number of times symbol j occurs in a sample of size n. The sequence $(N_n^j)_{j\geq 1}$ is called the *type* of the sample. Let $K_{n,r}$ denote the *occupancy counts*, defined as the number of symbols that appear exactly r times in a sample of size n:

$$K_{n,r} = \sum_{j\geq 1} \mathbb{I}_{\{N_n^j = r\}} \,.$$

The sequence $(K_{n,r})_{r\geq 1}$ is called the *profile* of the sample. The occupancy counts combine to yield the total number of distinct symbols in the sample, denoted by K_n:

$$K_n = \sum_{j\geq 1} \mathbb{I}_{\{N_n^j > 0\}} = \sum_{r\geq 1} K_{n,r} \,.$$

We will also encounter a quantity known as the *missing mass*, which corresponds to the cumulated probability of the unseen symbols:

$$M_{n,0} = \sum_{j\geq 1} p_j \mathbb{I}_{\{N_n^j = 0\}} \,.$$

Those quantities play a central role in estimating various properties of the sampling distribution, especially in situations where the sample size is small compared to the alphabet size [34, 45, 28]. In light of the correspondence between compression and probability estimation, they also are crucial to the analysis of the coding strategy we investigate here.

Henceforth, $\mathbb{E}$ will denote expectation with respect to the source distribution P or its n-fold product P^n, while $\mathbb{E}_f$ will denote the expectation with respect to the envelope distribution $(f_j)_{j\geq 1}$ or its n-fold product.

B. Karlin's framework

In 1967, S. Karlin introduced an elegant framework for investigating the asymptotic properties of infinite urn models. Sources over countable alphabets fit nicely into this framework. This framework has recently received attention in random combinatorics

and stochastic processes theory [See 24, for a survey]. Although the framework had been introduced in order to perform a thorough exploration of asymptotic normality of occupancy counts, it has also proved convenient when assessing the tightness of non-asymptotic tail or moment bounds for occupancy counts as demonstrated in [4].

A probability mass function $(p_j)_{j\geq 1}$ defines a *counting measure* ν defined by $\nu(\mathrm{d}x) = \sum_{j\geq 1} \delta_{p_j}(\mathrm{d}x)$, where δ_p is the Dirac mass at p. Let the counting function $\vec{\nu}(\cdot)$ be the right tail of ν, that is for all $x \in (0,1]$,

$$\vec{\nu}(x) = \nu[x,\infty[= |\{j \geq 1,\, p_j \geq x\}|\,.$$

Let us also define the measure ν_1 by

$$\nu_1(\mathrm{d}x) = \sum_{j\geq 1} p_j \delta_{p_j}(\mathrm{d}x)\,.$$

For $x \in [0,1]$, $\nu_1[0,x] = \sum_{j\geq 1} p_j \mathbb{I}_{p_j \leq x}$ is the cumulated probability of symbols with probability smaller than x. Expected occupancy counts and masses can be written simply as integrals against the measure ν. For instance

$$\mathbb{E}K_n = \int_0^1 (1-(1-x)^n)\,\nu(\mathrm{d}x)\,.$$

The sequence $(\mathbb{E}K_n)_n$ characterizes the measure ν (this follows easily from a Laplace transform argument).

C. Regular variation and regularly varying envelopes

The envelope distribution considered in this paper will most often satisfy a so-called *regular variation* property. A fascinating treatment of regular variation can be found in the textbook [5].

Definition 2. A measurable function $g :]a,\infty) \to]0,\infty)$ (where $a > 0$) is said to be regularly varying at infinity if for all $x > 0$

$$\lim_{t\to\infty} \frac{g(tx)}{g(t)} = x^\alpha$$

for some $\alpha \in \mathbb{R}$. This is summarized by $g \in \mathrm{RV}_\alpha$. If $\alpha = 0$, g is said to be slowly varying. Note that $g \in \mathrm{RV}_\alpha$ if and only if, for all $t > a$, $g(t) = t^\alpha \ell(t)$ where ℓ is slowly varying.

Useful results from regular variation theory are gathered in Appendix A.

Following Karlin [29] and Gnedin, Hansen, and Pitman [24], source classes with regularly varying envelopes are defined as follows.

Definition 3 (REGULARLY VARYING SOURCE CLASSES)**.** The envelope class $\mathcal{C}(f)$ is said to be regularly varying with index $\alpha \in [0,1]$ if the function $\vec{\nu}_f(1/\cdot)$ is regularly varying with index α (denoted $\vec{\nu}_f(1/\cdot) \in \mathrm{RV}_\alpha$).

Note that the property pertains to the envelope counting function. The counting functions associated with source distributions in the class may or may not satisfy the regular variation property.

If the envelope distribution is a so-called power-law distribution, that is if, for some $C > 0$ and $0 < \alpha < 1$, $f(j) \sim Cj^{-1/\alpha}, j \to \infty$, then $\vec{\nu}_f$ is regularly varying with index α [See 1, for details].

The case $\alpha = 0$ corresponds to light-tailed envelopes. It contains for instance frequencies proportional to $C\exp\left(-C' j^\beta\right)$, for some $\beta > 0$ (in particular, when $\beta = 1$, the Geometric distribution), or for example frequencies corresponding to the integral part of a log-normal variable as treated in [7].

Under the regular variation assumption, the asymptotic expected values of K_n, $K_{n,r}$ and $M_{n,0}$ are nicely related to the counting function $\vec{\nu}$ (see Appendix B). Appendix C relates properties of the counting function and properties of the quantile function of the sampling distribution.

If f, g denote two functions from $\mathbb{R}_+$ to $\mathbb{R}_+$, notation $f \asymp g$ means that there exists some constant $C \geq 1$ and $t_0 \in \mathbb{R}_+$ such that for $t > t_0$, $1/Cg(t) \leq f(t) \leq Cg(t)$.

III. Main results

We first describe the minimax redundancy of envelope classes $\mathcal{C}(f)^n$ in Karlin's framework, that is using the counting function $\vec{\nu}$ and measure ν_1. The first theorem follows from the Poissonization approach promoted in [49] and [1] [See also 27]. It gathers upper and lower bounds on the minimax redundancy of envelope classes. Note that the bounds do not depend on any regular variation assumption on the envelope function.

Theorem 4. *For any envelope function f, the minimax redundancy of $\mathcal{C}(f)^n$ satisfies*

i)

$$\overline{R}\left(\mathcal{C}(f)^n\right) \leq \log(\mathrm{e})\left(\int_1^n \frac{\vec{\nu}_f(1/t)}{2t}\mathrm{d}t + \vec{\nu}_f(1/n) + n\nu_{1,f}[0,1/n]\right) + O(\ell_f \log n)\,. \tag{2}$$

ii) Let $m = n - n^{2/3}$. For any envelope function f, there exists a constant $c_f > 0$ such, for large enough n,

$$\overline{R}\left(\mathcal{C}(f)^n\right) \geq \left(\log(\mathrm{e}) \int_1^m \frac{\vec{\nu}_f(1/t)}{2t}\mathrm{d}t - 5\vec{\nu}_f(1/m) - 1\right) \vee \left(\mathbb{E}_f[K_n] - c_f\right) . \tag{3}$$

Note that, as soon as the support of $(f_j)_{j\geq 1}$ is infinite, $\overline{R}\left(\mathcal{C}(f)^n\right) \gg \log n$. Hence, in our context, the term $O(\ell_f \log n)$ in (2) is asymptotically negligible with respect to the other terms.

For an envelope function f defined by frequencies $(f_j)_{j\geq 1}$, with counting function $\vec{\nu}_f$, recall that $\mathsf{R}_f(n)$ is defined as

$$\mathsf{R}_f(n) = \log(\mathrm{e}) \int_1^n \frac{\vec{\nu}_f(1/t)}{2t}\mathrm{d}t\,.$$

Regular variation assumptions on the envelope allow us to compare the different terms in (2) and (3) (see Theorems 19 and 23 in Appendix A and B). Assuming that the counting function $\vec{\nu}_f$ satisfies a regular variation property, the sequence $\left(\mathsf{R}_f(n)\right)_n$ is proportional to the redundancy rate of the envelope class $\mathcal{C}(f)$. It will also turn out to characterize the maximum redundancy of the PC code over $\mathcal{C}(f)$.

When $\vec{\nu}_f(1/\cdot) \in \mathrm{RV}_0$, Karamata's integration Theorem and Proposition 15 in [24] imply

$$\mathsf{R}_f(n) \gg \vec{\nu}_f(1/n) \underset{n\to\infty}{\sim} \mathbb{E}_f[K_n] \gg n\nu_{1,f}[0, 1/n]\,.$$

Hence if $\vec{\nu}_f(1/\cdot) \in \mathrm{RV}_0$, Theorem 4 entails $\overline{R}(\mathcal{C}(f)^n) \underset{n\to\infty}{\sim} \mathsf{R}_f(n)$.

When $\vec{\nu}_f(1/\cdot) \in \mathrm{RV}_1$, the largest term in the upper bound is now $n\vec{\nu}_{1,f}[0, 1/n]$, which is of order $\mathbb{E}_f[K_n] \gg \vec{\nu}_f(1/n)$.

Last but not least, when $\vec{\nu}_f(1/\cdot) \in \mathrm{RV}_\alpha$ for $0 < \alpha < 1$, all three terms in the right-hand-side of (2) have the same order of magnitude $\mathbb{E}_f[K_n] \asymp \vec{\nu}_f(1/n) \asymp \mathsf{R}_f(n)$. Hence, the following theorem follows directly from Theorem 4.

Theorem 5 ([1])**.** *Let $\mathcal{C}(f)$ be an envelope source class, with $f \in \mathrm{RV}_\alpha$ and $\alpha \in [0, 1[$. Then*

$$\overline{R}(\mathcal{C}(f)^n) \asymp \mathsf{R}_f(n)\,.$$

If $\alpha = 0$,

$$\lim_{n\to\infty} \frac{\mathsf{R}_f(n)}{\overline{R}(\mathcal{C}(f)^n)} = 1\,.$$

If $\alpha = 1$,

$$\overline{R}(\mathcal{C}(f)^n) \asymp \mathbb{E}_f[K_n] \gg \mathsf{R}_f(n)\,.$$

The implementation of the Poissonization technique by Acharya et al. answered a number of open questions raised in [8, 7]. Previously, the minimax redundancy rate for source classes defined by regularly varying envelopes was only known when envelopes were discrete log-concave (which implies non-decreasing hazard rate), that is for a strict subset of classes defined by slowly varying envelopes. Before [1], the best known general upper bound on the redundancy of envelopes classes was stated using the tail envelope function and read as

$$\overline{R}(\mathcal{C}(f)^n) \leq \inf_{u\leq n} \left\{ n\overline{F}(u)\log \mathrm{e} + \frac{u-1}{2}\log n \right\} + 2$$

where $\overline{F}(u) = \sum_{j>u} f_j$ [8].

Optimizing over u leads to outline $u = \vec{\nu}_f(\frac{\ln n}{2n})$. The redundancy upper bound then translates into

$$\overline{R}(\mathcal{C}(f)^n) \leq \left(\frac{2n}{\ln n}\nu_{1,f}\left(0, \tfrac{\ln n}{2n}\right) + \vec{\nu}_f\left(\frac{\ln n}{2n}\right)\right) \frac{\log n}{2}\,. \tag{4}$$

When $\vec{\nu}_f(1/t) \underset{t\to\infty}{\sim} t^\alpha \ell(t)$ with $\ell \in \mathrm{RV}_0, \alpha \in (0, 1)$, combining (4) and Proposition 13 from [24] reveals that the upper bound in (4) is asymptotically equivalent to

$$\frac{\log \mathrm{e}}{2^{1-\alpha}(1-\alpha)} \ell\left(\tfrac{2n}{\ln n}\right) n^\alpha (\ln n)^{1-\alpha}\,.$$

Invoking Proposition 13 from [24] and Karamata's integration Theorem show that the right-hand-side of (2) is asymptotically equivalent to

$$\frac{\log(\mathrm{e})(1+\alpha)}{2\alpha(1-\alpha)} \ell(n) n^\alpha\,.$$

For envelopes with positive regular variation index α, the redundancy upper bound from [8] is off by a factor $(\ln n)^{1-\alpha}$.

For envelopes that satisfy the discrete log-concavity condition (a special case of envelope function with non-decreasing hazard rate as treated in [7]), then by Lemma 24, the redundancy bounds derived in [7] and $\mathsf{R}_f(n)$ are asymptotically equivalent.

An interesting benefit from the Poissonization method concerns slowly varying envelopes that do not satisfy the discrete log-concavity condition. For slowly varying envelopes, the upper bound in (4) is asymptotically equivalent to $\vec{\nu}_f\left(\frac{\ln n}{2n}\right)\frac{\log n}{2}$. If we focus on $\vec{\nu}_f(1/t)=\lfloor\exp\left((\ln t)^\beta\right)\rfloor$ for $\beta\in(0,1), t>1$, then, integration by part shows that $\mathsf{R}_f(n)\le\vec{\nu}_f(1/n)(\ln(n))^{1-\beta}/\beta$. Under this setting, for $\beta\in(0,1/2)$, by Theorem from [5], $\vec{\nu}_f\left(\frac{\ln n}{2n}\right)\underset{n\to\infty}{\sim}\vec{\nu}_f(1/n)$. This is enough to conclude that the redundancy upper bound derived from (4) and the bound derived from (2) are not of the same order of magnitude: $\mathsf{R}_f(n)=o\big(\vec{\nu}_f\left(\frac{\ln n}{2n}\right)\frac{\log n}{2}\big)$.

We now turn to adaptivity and characterize the performance of the PC Code $(Q_n)_n$ defined in Section V on regularly varying envelope classes $\mathcal{C}(f)$, with index $\alpha\in[0,1[$.

A regular variation assumption on the envelope distribution is enough to characterize the redundancy rate of the source class (Theorem 5). Underneath this rewarding message lies a more ambivalent finding: the redundancy rate breaks down into three components whose relative importance depends on the unknown regular variation index (Theorem 4).

Experience with different attempts at designing adaptive codes suggests that an adaptive code should be able to sort rare symbols from frequent symbols (with probability mass larger than the reciprocal of the current message length) and to work with an effective alphabet of size not larger than $\vec{\nu}_f(1/n)$. The PC code attempts to do this in a simple and efficient manner: any recurring symbol is deemed to be frequent; any symbol occurring for the first time is deemed to be rare. In this way, the size of the alphabet used by the mixture coder coincides with the number of distinct symbols in the portion of the message encoded so far. The latter quantity is well concentrated around its mean value and is bounded from above by the mean value under the envelope distribution. No special attempt is made to encode the dictionary in a way that adapts to a regular variation index since the source distribution may not satisfy any regular variation assumption.

Despite its simplicity the PC code achieves adaptivity within a $\log\log n$ factor over all regularly varying envelope classes.

Theorem 6. *Let $(Q_n)_n$ be the sequence of coding distributions associated with the Pattern Censoring code. For all $\alpha\in[0,1[$, for every envelope function f with $\vec{\nu}_f(1/\cdot)\in\mathrm{RV}_\alpha$, there exists constants $a_f,b_f>0$ such that*

$$(a_f+o_f(1))\le\frac{\overline{R}(\mathcal{C}(f)^n)}{\mathsf{R}_f(n)}\le\frac{\overline{R}(Q_n,\mathcal{C}(f)^n)}{\mathsf{R}_f(n)}\le(b_f+o_f(1))\log\log n\,.$$

In particular, the Pattern Censoring code is adaptive, within a $\log\log n$ factor, with respect to the collection

$$\left\{\mathcal{C}(f):f\in\mathrm{RV}_\alpha,\alpha\in[0,1[\right\}.$$

The PC code is designed in a way that parallels the design of the AC and ETAC codes from [7, 10]. All three codes interleave mixture coding of symbols that are deemed frequent and Elias encoding of symbols that are deemed rare. They differ in the way censoring is performed. The AC code censors records that is symbols that are larger than symbols seen so far. The ETAC code censors the n^{th} symbol if it is larger than M_n where M_n is a carefully chosen order statistics of the sub-sample $X_1,X_2,\ldots,X_n$ (See Equation (13) in Appendix C). At each instant, the AC and the ETAC codes handle an effective alphabet formed by a sequence of consecutive symbols. The AC code is doomed to failure if the source has a positive regular variation index $\alpha\in(0,1)$: the increasing alphabet used by the mixture code typically grows like $n^{\alpha/(1-\alpha)}$ instead of n^α. The increasing alphabet used by the mixture code of the ETAC code grows much more reasonably: if the source distribution is dominated by a regularly varying envelope f with index $\alpha\in(0,1)$, the increasing alphabet typically grows like $\mathbb{E}M_n$ which is or order $\vec{\nu}_f(1/n)$ (See Appendix C). If the envelope is slowly varying and discrete log-concave, $\mathbb{E}M_n$ also grows at a pace that is consistent with the slow variation property of the envelope. But our current understanding of the rate of slow variation problem [See 5, Section 2.3] does not allow us to quantify precisely the rate of growth the alphabet when the envelope is slowly varying and is equivalent to a function in the de Haan class (See Definition 22 in Appendix A) where the auxiliary function tends to infinity. This lack of understanding also prevents us from quantifying the redundancy of the ETAC code with respect to corresponding envelope classes.

The ETAC and PC codes also differ in the way they encode the escape symbol 0 (see below for details). The mixture encoder used by ETAC code always considers 0 as a new symbol: the predictive probability associated with 0 at instant $n+1$ is $1/(n+(M_n+1)/2)$. The predictive probability assigned to 0 should rather be close to an estimator of the probability of discovering a new symbol, that is of the missing mass. For a regularly varying sampling distribution with index $\alpha\in(0,1)$, the predictive probability assigned to 0 should scale like $n^{\alpha-1}$ rather than $1/n$. The PC code assigns a predictive probability $(K_n+1/2)/(n+K_n+1)$ to the escape symbol. Even though this does not coincide with the Good-Turing estimator of the missing mass $K_{n,1}/n$, it scales correctly with n for regularly varying sampling distributions.

Theorems 4 and 5 are established in Section IV. Theorem 6 is proved in Section VI. Useful technical results and arguments are stated in Section VII.

IV. Minimax redundancy of envelope source classes

This section describes upper and lower bounds for the minimax redundancy of envelope source classes. The techniques are borrowed from [1] where tight bounds on *minimax regret* (see below) are derived thanks to Poissonization arguments. Not too surprisingly, Poissonization allows to derive tight bounds for minimax redundancy as well. In Karlin's framework these bounds

are illuminating. When applied to regularly varying envelopes with positive regular variation parameter, the resulting bounds are tight up to constant factors. When applied to slowly varying envelopes, the resulting bounds actually imply asymptotic equivalence.

We now give the main ingredients leading to Theorem 4. The upper bound (2) follows from minimax regret bounds in [1]. Upper bounds are then translated using Karlin's framework, this provides us with insights onto effective coding methods. In the other direction, we need to show that Poissonization arguments also apply to minimax redundancy, and to establish lower bounds on the redundancy of Poisson classes.

A. Properties of minimax redundancy

The next theorem collects properties of minimax redundancies [See 23, for a general perspective].

Proposition 7. *The sequence $(\overline{R}(\mathcal{C}^n))_{n\geq 1}$ of minimax redundancies of a class $\mathcal{C}$ of stationary memoryless sources satisfies*

i) $(\overline{R}(\mathcal{C}^n))_{n\geq 1}$ is non-decreasing.

ii) $(\overline{R}(\mathcal{C}^n))_{n\geq 1}$ is sub-additive:

$$\text{for all} \quad n,m\geq 1, \quad \overline{R}(\mathcal{C}^{n+m}) \leq \overline{R}(\mathcal{C}^n) + \overline{R}(\mathcal{C}^m)\,.$$

iii) The minimax redundancy is equal to the maximin Bayesian redundancy:

$$\begin{aligned}\overline{R}(\mathcal{C}^n) &= \sup_{\pi\in\mathcal{M}_1(\mathcal{C})} \inf_{Q_n\in\mathcal{M}_1(\mathcal{X}^n)} \int D(P^n,Q_n)\mathrm{d}\pi(P)\\ &= \sup_{\pi\in\mathcal{M}_1(\mathcal{C})} \int D(P^n,P_\pi^n)\mathrm{d}\pi(P)\,,\end{aligned}$$

where P_π^n is the mixture distribution given by $P_\pi(j)=\int P(j)\mathrm{d}\pi(P)$.

We will also use the notion of regret that originates from individual sequence analysis [12]. The regret of a coding probability with respect to a source class $\mathcal{C}$ is

$$\widehat{R}(Q_n,\mathcal{C}^n) := \sup_{x_{1:n}\in\mathcal{X}^n} \sup_{P\in\mathcal{C}} \log\frac{P^n(x_{1:n})}{Q_n(x_{1:n})}$$

and the minimax regret of $\mathcal{C}^n$ is $\widehat{R}(\mathcal{C}^n) := \inf_{Q_n}\widehat{R}(Q_n,\mathcal{C}^n)$.

As pointed out by Acharya et al. [1], the minimax redundancy of a class of stationary memoryless sources is equal to the minimax redundancy of the induced class on *types*. Recall that the type of a sequence [16] is the sequence of frequencies of the different symbols in the sequence. As such, it is a random variables with values in the set of infinite sequences of integers with finitely many non-zero entries. More precisely, for $P\in\mathcal{C}$, let us denote by $\tau(P^n)$ the distribution of the *type* of a sequence $(X_1,\ldots,X_n)\sim P^n$, that is $\tau(P^n)$ is the probability distribution of the sequence $(N_n^j)_{j\geq 1}$. For a class $\mathcal{C}$ of sources over $\mathcal{X}$, the class of probability distributions $\tau(\mathcal{C}^n)$ is defined as as

$$\tau(\mathcal{C}^n) = \{\tau(P^n),\, P\in\mathcal{C}\}\,.$$

As types are minimal sufficient statistics,

$$\overline{R}(\mathcal{C}^n) = \overline{R}(\tau(\mathcal{C}^n)) \quad \text{and} \quad \widehat{R}(\mathcal{C}^n) = \widehat{R}(\tau(\mathcal{C}^n))\,.$$

Poissonization provides a simplified framework in various settings. Under *Poisson sampling*, the message length is a random variable N that is distributed according to a Poisson distribution with mean n ($N\sim\mathcal{P}(n)$). Picking a Poissonized sample from P consist first of picking N according to $\mathcal{P}(n)$ and then picking an i.i.d. sample of random length N from P.

Let $\mathcal{C}^{\mathcal{P}(n)}$ be the Poissonized version of $\mathcal{C}^n$:

$$\mathcal{C}^{\mathcal{P}(n)} = \left\{P^{\mathcal{P}(n)},\, P\in\mathcal{C}\right\},$$

where, for all $P\in\mathcal{C}$, all $k\geq 1$ and $x_{1:k}\in\mathcal{X}^k$,

$$P^{\mathcal{P}(n)}(x_{1:k}) = \mathbb{P}(N=k).P^k(x_{1:k}).$$

By convention, the empty sequence has probability zero. For $j\geq 1$, let us denote by $N_j(n)$ the number of occurrences of symbol j in a Poisson sample with size $\mathcal{P}(n)$. Then $N_j(n)$ is distributed as $\mathcal{P}(nP(j))$, and a very useful property of Poisson sampling is that the symbol counts $(N_j(n))_{j\geq 1}$ are independent. Let us also note that, as in the fixed-n setting, the redundancy of $\mathcal{C}^{\mathcal{P}(n)}$ is equal to the type-redundancy [1]:

$$\overline{R}(\mathcal{C}^{\mathcal{P}(n)}) = \overline{R}\left(\tau(\mathcal{C}^{\mathcal{P}(n)})\right) \quad \text{and} \quad \widehat{R}(\mathcal{C}^{\mathcal{P}(n)}) = \widehat{R}\left(\tau(\mathcal{C}^{\mathcal{P}(n)})\right)\,. \tag{5}$$

Proposition 8. *For all class $\mathcal{C}$, the Poissonized minimax redundancy satisfies*

$$\overline{R}(\mathcal{C}^{\mathcal{P}(n)}) = \inf_{(Q_k)} \sup_{P\in\mathcal{C}} \sum_{k\geq 0} \mathbb{P}(N=k) D(P^k, Q_k) = \sum_{k\geq 0} \mathbb{P}(N=k)\overline{R}(\mathcal{C}^k)\,,$$

where the infimum is taken over sequences $(Q_k)_{k\geq 0}$, such that, for all $k\geq 0$, Q_k is a probability distribution over $\mathcal{X}^k$.

Minimax regret and minimax regret under Poisson sampling are related by Theorem 2 from [1]. Falahatgar et al. [20] also show that $\overline{R}(\mathcal{C}^{\mathcal{P}(n)}) \leq 2\overline{R}(\mathcal{C}^n)$. The next proposition complements these connections.

Proposition 9. *For any class $\mathcal{C}$ of stationary memoryless sources with $\overline{R}(\mathcal{C}) < \infty$,*

$$\overline{R}\left(\mathcal{C}^{\mathcal{P}(n-n^{2/3})}\right) + o_{\mathcal{C}}(1) \leq \overline{R}(\mathcal{C}^n) \leq (1-o(1))\overline{R}\left(\mathcal{C}^{\mathcal{P}(n+n^{2/3})}\right)$$

B. Minimax redundancy of envelope classes

When the source class is an envelope source class $\mathcal{C} = \mathcal{C}(f)$, the class of types distributions $\tau\left(\mathcal{C}(f)^{\mathcal{P}(n)}\right)$ under Poisson sampling is a class of product distributions,

$$\tau\left(\mathcal{C}(f)^{\mathcal{P}(n)}\right) = \left\{ \prod_{j=1}^{\infty} \mathcal{P}(np_j) : P = (p_j)_{j\geq 1} \in \mathcal{C}(f) \right\}. \tag{6}$$

Let us define, for $\lambda \geq 0$,

$$\mathcal{P}^\star(\lambda) = \{\mathcal{P}(\mu) : \mu \leq \lambda\}\,,$$

the class of Poisson distributions with mean smaller than λ. The minimax redundancy of envelope classes under Poisson sampling is tightly related to the sum of minimax redundancies of classes of Poisson distributions.

Lemma 10. *Let $\mathcal{C}(f)$ be an envelope class. Then for all $n\geq 0$,*

$$\sum_{j=\ell_f}^{\infty} \overline{R}(\mathcal{P}^\star(nf_j)) \leq \overline{R}\left(\mathcal{C}(f)^{\mathcal{P}(n)}\right) \leq \ell_f \overline{R}(\mathcal{P}^\star(n)) + \sum_{j=\ell_f}^{\infty} \overline{R}(\mathcal{P}^\star(nf_j))\,.$$

The next lemma provides us with a lower bound on $\overline{R}(\mathcal{P}^\star(nf_j))$.

Lemma 11. *For $\lambda \geq 1$, the redundancy of $\mathcal{P}^\star(\lambda)$ satisfies*

$$\overline{R}(\mathcal{P}^\star(\lambda)) \geq \frac{\log\lambda}{2} - 5\,.$$

The first lower bound in Theorem 4 now follows from Proposition 9, Lemma 10, and Lemma 11.

V. The Pattern Censoring code

The Pattern Censoring code (PC code) borrows ideas from pattern coding [32, 22]. The encoder and the decoder maintain a dictionary of symbols seen so far. With each symbol is associated the number of occurrences of the symbol and its rank of insertion in the dictionary. If the n^{th} symbol in the message is the first occurrence of symbol j, an escape symbol is fed to the Krichevsky-Trofimov encoder, symbol j is inserted into a dictionary with count 1 and symbol j is encoded using a general purpose prefix code for integer (an Elias encoder). If the n^{th} symbol has already been inserted into the dictionary, the index of its first occurrence is fed to the Krichevsky-Trofimov encoder.

We first give a bird-eye view of the PC encoding of a sequence $x_{1:n} = x_1, \ldots, x_n$ of symbols from the message alphabet $\mathcal{X} = \mathbb{N}_+ = \mathbb{N}\setminus\{0\}$. The message is appended with a 0 that will serve as a termination signal.

Encoding:

— The dictionary is initialized as follows: $\mathcal{D}_0 = \{\langle 0, 0\rangle\}$, *i.e.* we start with the single (escape) symbol 0, associated with its virtual *rank of insertion*, 0.

— At every index i corresponding to an input symbol, maintain a dictionary $\mathcal{D}_i$ containing symbol 0 and the symbols occurring in $x_{1:i}$, along with their rank of insertion. Dictionary $\mathcal{D}_i$ thus contains $\langle 0, 0\rangle$ and ordered pairs $\langle j, k\rangle$ if symbol j is the k^{th} distinct symbol occurring in $x_{1:i}$ ($\exists m \leq i, x_m = j, K_{m-1} = k-1, K_m = k$).

— Create a *censored* sequence $\widetilde{x}_{1:n}$ such that every symbol x_i that does not belong to $\mathcal{D}_{i-1}$ is replaced by the special 0 symbol and every symbol x_i that belongs to $\mathcal{D}_{i-1}$ is replaced by its rank of insertion:

$$\widetilde{x}_i = \begin{cases} k & \text{if } \exists k \leq K_{i-1}, \quad \langle x_i, k\rangle \in \mathcal{D}_{i-1} \\ 0 & \text{otherwise}\,. \end{cases}$$

Note that each time i such that $\widetilde{x}_i = 0$ corresponds to the discovery of a new symbol and thus to an insertion in the dictionary: $\mathcal{D}_i = \mathcal{D}_{i-1} \cup \langle x_i, K_i \rangle$, with $K_i = K_{i-1} + 1$.

— Let K_n be the number of *redacted* (censored-out) input symbols (this is the number of distinct symbols in $x_{1:n}$) and let $i_{1:K_n}$ be their indexes. Extract the subsequence $x_{i_{1:K_n}}$ of all such redacted symbols.
— Perform an instantaneously decodable lossless progressive encoding (in the style of Mixture / arithmetic coding) of the censored sequence $\widetilde{x}_{1:n}$, assuming decoder side-information about past symbols. Call the resulting string C_M.
— Perform an instantaneously decodable lossless encoding (in the style of Elias / integer coding) of each redacted symbol in $x_{i_{1:K_n}}$ and of the appended 0 individually rather than as a sequence, assuming decoder side-information about past symbols. Call the resulting string C_E.
— Interleave the coded redacted symbols of C_E just after each coded 0 symbol of C_M, to form the overall PC-code.
— The first 0 in the sequence of redacted symbols signals termination.

Decoding:

— Decode the interleaved C_M and C_E strings until exhaustion, as follows.
— Decode C_M to obtain $\widetilde{x}_{1:n}$ progressively.
— When a 0 is encountered, decode the single interleaved redacted symbol from C_E to take the place of the 0 symbol in the decoded sequence, then move back to decoding C_M.
— Note that at all times the decoder knows the entire past sequence, and therefore the decoder side past side-information hypothesis is upheld.

The censored sequence generated by message *abracadabra* is

$$00010101231$$

while the dictionary constructed from this sequence is

$$\{\langle 0, 0\rangle, \langle a, 1\rangle, \langle b, 2\rangle, \langle r, 3\rangle, \langle c, 4\rangle, \langle d, 5\rangle\}$$

We now give the details of the dictionary, the encoding of the censored sequence $\widetilde{x}_{1:n}$, and the encoding of the redacted symbols $x_{i_{1:K_n}}$. We also take additional care in guaranteeing that the overall PC code is instantaneously decodable.

The censored sequence $\widetilde{x}_{1:n}$ is encoded into the string C_M as follows. We start by appending an extra 0 at the end of the original censored sequence, to signal the termination of the input. We therefore in fact encode $\widetilde{x}_{1:n}0$ into C_M. We do this by performing a progressive arithmetic coding [36] using coding probabilities $\widetilde{Q}_{n+1}(\widetilde{x}_{1:n}0)$ given by:

$$\widetilde{Q}_{n+1}(\widetilde{x}_{1:n}0) = \widetilde{Q}_{n+1}(0 \mid \widetilde{x}_{1:n}) \prod_{i=0}^{n-1} \widetilde{Q}_{i+1}(\widetilde{x}_{i+1} \mid \widetilde{x}_{1:i}) \,,$$

where the predictive probabilities $\widetilde{Q}_{i+1}(\cdot \mid \widetilde{x}_{1:i})$ are given by Krichevsky-Trofimov mixtures on the alphabet $\{0, 1, \ldots, K_i\}$,

$$\widetilde{Q}_{i+1}\left(\widetilde{X}_{i+1} = k \mid \widetilde{X}_{1:i} = \widetilde{x}_{1:i}\right) = \frac{\widetilde{n}_i^k + \frac{1}{2}}{i + \frac{K_i+1}{2}} \,,$$

where, for $0 \le k \le K_i$, $\widetilde{n}_i^k$ is the number of occurrences of symbol k in $\widetilde{x}_{1:i}$. Note that, for $0 \le k \le K_i$,

$$\widetilde{n}_i^k = \begin{cases} K_i & \text{if } k = 0 \ , \\ n_i^j - 1 & \text{if } 1 \le k \le K_i \text{ and } \langle j, k\rangle \in \mathcal{D}_i \ , \end{cases}$$

where n_i^j is the number of occurrences of symbol j in $x_{1:i}$.

The subsequence $x_{i_{1:K_n}}$ of redacted symbols is encoded into the string C_E as follows. For each $i \in i_{1:K_n}$, we encode x_i using Elias penultimate coding [19]. Thus, if $x_i = j$ and $N_{i-1}^j = 0$, the cost of inserting this new symbol in the dictionary (the Elias encoding of j) is $1 + \log(j \vee 1) + 2\log(1 + \log(j \vee 1))$. The extra 0 appended to the message is fed to the Elias encoder. Since no other encoded redacted symbol but this one equals 0, it unambiguously signals to the decoder that the last 0 symbol decoded from C_M is in fact the termination signal. This ensures that the overall code is instantaneously decodable, and that it corresponds to a coding probability Q_n.

The number of elementary operations required by dictionary maintenance depends on the chosen implementation [See 43]. If a self-adjusting balanced search tree is chosen, the number of elementary operations is at most proportional to the logarithm of the size of the dictionary. As the expected total number of symbols inserted in the dictionary coincides with the number of distinct symbols K_n, on average it is sub-linear, the total expected computational cost of dictionary maintenance is $O(n \log n)$ if n denotes the message length.

After reading the i^{th} symbol from the message, the alphabet used by the Krichevsky-Trofimov encoder is $0, \ldots, K_i$, the state of the arithmetic encoder is a function of the counts $\widetilde{n}_i^0, \widetilde{n}_i^1, \ldots \widetilde{n}_i^{K_i}$. Counts may be handled using map or dictionary data structures provided by modern programming languages.

Algorithm 1 PC encoder

Append 0 at the end of message $x_{1:n}$
$\mathcal{D} \leftarrow \{\langle 0, 0\rangle\}$ {*Initialize dictionary*}
$n^0 \leftarrow 0$ {*Initialize counters*}
Initialize the arithmetic encoder
$K \leftarrow 0$ {*Initialize the number of distinct symbols seen so far*}
for $i = 1$ **to** length$(x_{1:n}0)$ **do**
  $j \leftarrow x_i$
  if symbol $j \in \mathcal{D}$ and $j \neq 0$ **then**
    $k \leftarrow$ rank of insertion of j in $\mathcal{D}$ { *exists* $\langle j, k\rangle \in \mathcal{D}$ }
    feed mixture arithmetic encoder with k
    emit output of arithmetic encoder if any
    $n^k \leftarrow n^k + 1$ {*update counter*}
  else
    feed mixture arithmetic encoder with escape symbol 0
    { *this forces the arithmetic coder to output the encoding of the current substring* }
    emit the output of the arithmetic encoder
    feed the Elias encoder with j
    emit the output of the Elias encoder
    $n^0 \leftarrow n^0 + 1$
    if $j \neq 0$ **then**
      $K \leftarrow K + 1$ {*increment number of distinct symbols*}
      $\mathcal{D} \leftarrow \mathcal{D} \cup \{\langle j, K\rangle\}$
      $n^K \leftarrow 1$ {*initialize new counter*}
    else
      exit
    end if
  end if
end for

From a bird-eye viewpoint, the sequence of censored symbols defines a parsing of the message into substrings of uncensored symbols that are terminated by 0. Each substring is encoded by an arithmetic encoder provided with incrementally updated sequences of probability vectors. The arithmetic encodings of substrings are interleaved with Elias encodings of censored symbols.

Encoding and decoding are performed in an incremental way, even though arithmetic coding may require buffering [40].

VI. Redundancy of the Pattern Censoring Code

The redundancy of the PC code, that is the cumulative entropy risk of the coding probability with respect to the sampling probability is the sum of expected instantaneous redundancies.

$$D(P^n, Q_n) = \sum_{i=0}^{n-1} \mathbb{E}_{P^i}\left[\mathbb{E}_P\left[\log \frac{P(X_{i+1})}{Q_{i+1}(X_{i+1}|X_{1:i})}\Big| X_{1:i}\right]\right] .$$

For $i \geq 0$, the conditional expected instantaneous redundancy is given by

$$\mathbb{E}_P\left[\log \frac{P(X_{i+1})}{Q_{i+1}(X_{i+1}|X_{1:i})}\Big| X_{1:i}\right]$$
$$= \underbrace{\sum_{j\geq 1} p_j \mathbb{I}_{N_i^j>0} \log\left(\frac{p_j\left(i+\frac{K_i+1}{2}\right)}{N_i^j - \frac{1}{2}}\right)}_{\text{(I)}} + \underbrace{\sum_{j\geq 1} p_j \mathbb{I}_{N_i^j=0} \log\left(\frac{p_j\left(i+\frac{K_i+1}{2}\right)}{K_i + \frac{1}{2}}\right)}_{\text{(II)}} + \underbrace{\sum_{j\geq 1} p_j \mathbb{I}_{N_i^j=0}\left(1 + \log(j) + 2\log\left(\log(j)+1\right)\right)}_{\text{(III)}} .$$

Terms (I) and (II) correspond to the mixture encoding on the censored sequence while the last term corresponds to the Elias encoding of first occurrences of symbols.

Let us point out right away that, in contrast with the analyses from [6, 7, 10], the analysis of the PC code does not separate the contributions of the two interleaved codes. We proceed cautiously in order to take into account the fact that redundancy accounting has to match the (unknown) envelope tail behavior.

The conditional instantaneous redundancy may be rearranged as follows:

$$\mathbb{E}_P\left[\log\frac{P(X_{i+1})}{Q_{i+1}(X_{i+1}|X_{1:i})}\Big|X_{1:i}\right] \tag{7}$$
$$= \underbrace{\sum_{j\geq 1} p_j \log\left(ip_j\left(1+\frac{K_i+1}{2i}\right)\right)}_{A:=} + \underbrace{\sum_{j\geq 1} p_j \mathbb{I}_{N_i^j>0}\log\left(\frac{1}{N_i^j-\frac{1}{2}}\right)}_{B:=} + \underbrace{\sum_{j\geq 1} p_j \mathbb{I}_{N_i^j=0}\left(1+\log\left(\frac{j}{K_i+\frac{1}{2}}\right)\right)}_{C:=} + \underbrace{2\sum_{j\geq 1} p_j \mathbb{I}_{N_i^j=0}\log\left(\log(j)+1\right)}_{D:=}.$$

Note that in term C the $\log(j)$ contribution of Elias encoding is balanced by the $\log(K_i+1/2)$ coming from the encoding of the 0 symbols in the censored sequence. This turns out to be crucial and allows us to obtain adaptivity within a $\log\log n$ factor (which comes from the residual part of Elias code) instead of a $\log n$ factor. The encoding of the 0 symbols in the censored sequence is then pivotal, this is one of the main differences with the AC or ETAC codes: in those previous codes, zeros were actually not encoded, and the KT code was working with the counts stemming from the sequence $X_{1:i}$ instead of $\widetilde{X}_{1:i}$. Here, the KT code considers 0 as a proper symbol in the sequence $\widetilde{X}_{1:i}$, and this allows us to gain the term (II) in (7) at the price of decreasing the other counts N_i^j by one, which does not significantly affects the redundancy.

Bounding the redundancy of the PC code then consists in two steps. The instantaneous redundancy incurred by encoding the $(i+1)^{\text{th}}$ symbol code is first upper bounded by an expression involving $\mathbb{E}K_i$, the expected number of distinct symbols in a message of length i as well as other occupancy counts. This should not be regarded as surprising if we keep in mind that redundancy and regret of pattern coding essentially depend on the profile of the message [33, 32, 35], that is on the occupancy counts. This first upper bound is distribution-free, it holds for all stationary memoryless sources.

Proposition 12. *For any source P, the instantaneous redundancy of the* PC *code satisfies*

$$\mathbb{E}_P\left[\log\frac{P(X_{i+1})}{Q_{i+1}(X_{i+1}|X_{1:i})}\right] \leq \kappa\frac{\mathbb{E}K_i}{i} + \sum_{j\geq 1} p_j \mathbb{P}\{N_i^j=0\}\left(\log\left(\frac{j}{\mathbb{E}K_i}\right)+2\log\left(\log(j)+1\right)\right),$$

where $\kappa \leq 19$.

The second step assumes that the source probability mass function is dominated by a regularly varying envelope.

Proposition 13. *Let $P\in\mathcal{C}(f)$ and assume that $\vec{\nu}_f(1/\cdot)\in \mathrm{RV}_\alpha$ with $\alpha\in[0,1[$. Then there exists $i_0\in\mathbb{N}$ such that for all $i\geq i_0$,*

$$\mathbb{E}_P\left[\log\frac{P(X_{i+1})}{Q_{i+1}(X_{i+1}|X_{1:i})}\right] \leq c_\alpha\frac{\log\log(i)\vec{\nu}_f(1/i)}{i},$$

with $c_\alpha = 3\left(\Gamma(1-\alpha)+\frac{1}{1-\alpha}\right)$.

Before proceeding with the analysis of the code, we first state some useful comparisons between the expected occupancy counts, the number of distinct symbols in a random message, the missing mass and the measure ν_1, under the source and under the envelope. We start again with distribution-free statements, and proceed with the precise asymptotic results which are valid in the regular variation scenario. As our code is fundamentally related to occupancy counts and to the occurrences of new symbols, those will be very helpful to evaluate the contribution of each term to the redundancy.

Lemma 14. *For any stationary memoryless source defined by the counting function $\vec{\nu}$, for all $i\geq 1$,*

i) The expected missing mass, the expected number of singletons and the expected number of distinct symbols are connected:

$$\mathbb{E}M_{i,0} \leq \mathbb{E}\frac{K_{i,1}}{i} \leq \mathbb{E}\frac{K_i}{i},$$

ii)

$$\frac{\mathrm{e}-1}{\mathrm{e}}\vec{\nu}(1/i) \leq \mathbb{E}K_i \leq \vec{\nu}(1/i) + i\nu_1(0,1/i)$$

iii)

$$1 \leq \mathbb{E}K_i \times \mathbb{E}[1/K_i] \leq 3.$$

The largest value κ_1 of $\mathbb{E}[1/K_i]\mathbb{E}K_i$ over all $i\in\mathbb{N}$ and over all sampling distributions is yet to be determined. It is not smaller than $9/8$: consider a geometric distribution with success probability $2/3$.

The expected number of distinct symbols under the sampling distribution is related to the expected number of distinct symbols under the envelope distribution.

Lemma 15. *If $P\in\mathcal{C}(f)$, then $\vec{\nu}(1/i)\leq \ell_f + \vec{\nu}_f(1/i)$ and $\mathbb{E}K_i \leq \ell_f + \mathbb{E}_f[K_i]$.*

Finally, if the counting function $\vec{\nu}_f$ associated with the envelope is regularly varying in the sense of [29], the expected number of distinct symbols under the envelope distribution is simply connected with $\vec{\nu}_f$.

Lemma 16. *[29, 24] If $\vec{\nu}_f(1/\cdot) \in \mathrm{RV}_\alpha$ with $\alpha \in [0,1[$, then*

$$\mathbb{E}_f[K_n] \underset{n\to\infty}{\sim} \Gamma(1-\alpha)\vec{\nu}_f(1/n)\,.$$

For all $\varepsilon > 0$, there exists $n_0 \in \mathbb{N}$, such that for $n \geq n_0$,

$$n\,\nu_{1,f}[0,1/n] \leq \frac{\alpha+\varepsilon}{1-\alpha}\vec{\nu}_f(1/n)\,.$$

Notice that $\Gamma(1) = 1$ whereas $\lim_{\alpha\nearrow 1}\Gamma(1-\alpha) = \infty$. The heavier the envelope tail, the larger the ratio between $\mathbb{E}_f K_n$ and $\vec{\nu}_f(1/n)$.

We now proceed with the analysis of the redundancy of the PC code.

Proof of Theorem 6: Recall expansion (7) of the instantaneous redundancy incurred by encoding the $(i+1)^{\text{th}}$ symbol, given $X_{1:i}$:

$$\mathbb{E}_P\left[\log\frac{P(X_{i+1})}{Q_{i+1}(X_{i+1}|X_{1:i})}\Big|X_{1:i}\right]$$
$$= \underbrace{\sum_{j\geq 1} p_j \log\left(ip_j\left(1+\frac{K_i+1}{2i}\right)\right)}_{A:=} + \underbrace{\sum_{j\geq 1} p_j \mathbb{I}_{N_i^j>0}\log\left(\frac{1}{N_i^j-\frac{1}{2}}\right)}_{B:=} + \underbrace{\sum_{j\geq 1} p_j \mathbb{I}_{N_i^j=0}\left(1+\log\left(\frac{j}{K_i+\frac{1}{2}}\right)\right)}_{C:=} + \underbrace{2\sum_{j\geq 1} p_j \mathbb{I}_{N_i^j=0}\log\left(\log(j)+1\right)}_{D:=}\,.$$

Starting with term A, we use the fact that for all $x \geq 0$, $\log(1+x) \leq \log(\mathrm{e})x$ and obtain

$$A \leq \sum_{j\geq 1} p_j \log(ip_j) + \log(\mathrm{e})\frac{K_i+1}{2i}\,.$$

Averaging over the first i symbols $X_{1:i}$ leads to

$$\mathbb{E}[A] \leq \sum_{j\geq 1} p_j \log(ip_j) + \log(\mathrm{e})\frac{\mathbb{E}[K_i]+1}{2i}\,.$$

Moving on to term B, the conditional Jensen inequality entails

$$\begin{aligned}\mathbb{E}[B] &= \mathbb{E}\left[\sum_{j\geq 1} p_j \mathbb{I}_{N_i^j>0}\mathbb{E}\left[\log\frac{1}{N_i^j-\frac{1}{2}}\,\middle|\, N_i^j>0\right]\right]\\ &\leq \mathbb{E}\left[\sum_{j\geq 1} p_j \mathbb{I}_{N_i^j>0}\log\mathbb{E}\left[\frac{1}{N_i^j-\frac{1}{2}}\,\middle|\, N_i^j>0\right]\right]\,.\end{aligned}$$

Resorting to Lemma 28 in Appendix E,

$$\begin{aligned}\mathbb{E}[B] &\leq \mathbb{E}\left[\sum_{j\geq 1} p_j \mathbb{I}_{N_i^j>0}\log\left(\frac{1}{ip_j}\left(1+\frac{9}{ip_j}\right)\right)\right]\\ &\leq -\sum_{j\geq 1} p_j \mathbb{P}(N_i^j>0)\log(ip_j) + \log(\mathrm{e})\frac{9\mathbb{E}[K_i]}{i}\,.\end{aligned}$$

Hence,

$$\mathbb{E}[A+B] \leq \frac{1}{i}\sum_{j\geq 1} ip_j \mathrm{e}^{-ip_j}\log(ip_j) + \log(\mathrm{e})\frac{10\mathbb{E}[K_i]}{i}\,.$$

For $p_j \leq 1/i$ or equivalently $j > \vec{\nu}(1/i)$, $\log(ip_j) \leq 0$. As for all $x \geq 1$, $x\exp(-x)\log(x) \leq \log(\mathrm{e})/\mathrm{e}$,

$$\begin{aligned}\frac{1}{i}\sum_{j\geq 1} ip_j \mathrm{e}^{-ip_j}\log(ip_j) &\leq \frac{1}{i}\sum_{j\leq\vec{\nu}(1/i)} ip_j \mathrm{e}^{-ip_j}\log(ip_j)\\ &\leq \frac{\vec{\nu}(1/i)}{i}\frac{\log(\mathrm{e})}{\mathrm{e}} \;\leq\; \frac{\log(\mathrm{e})}{\mathrm{e}-1}\frac{\mathbb{E}K_i}{i}\,,\end{aligned}$$

where the last inequality comes from Lemma 14, ii). We get

$$\mathbb{E}[A+B] \leq \left(\frac{\log(\mathrm{e})}{\mathrm{e}-1} + 10\log(\mathrm{e})\right)\frac{\mathbb{E}[K_i]}{i}\,. \tag{8}$$

We turn towards term C. As $K_i + 1/2 \geq \frac{1}{2}(K_i+1)$, we have

$$\begin{aligned} C &\leq 2M_{i,0} + \sum_{j\geq 1} p_j \mathbb{I}_{N_i^j=0} \log\left(\frac{j}{K_i+1}\right) \\ &= 2M_{i,0} + \sum_{j\geq 1} p_j \mathbb{I}_{N_i^j=0} \log\left(\frac{j}{\sum_{\ell\neq j}\mathbb{I}_{N_i^\ell>0}+1}\right) . \end{aligned}$$

The collection of random variables $(N_i^j)_{j\geq 1}$ is *negatively associated* (See Appendix D). Hence, for all $j\geq 1$, N_i^j and $\sum_{\ell\neq j} N_j^\ell$ are negatively associated,

$$\mathbb{E}\left[\mathbb{I}_{N_i^j=0}\log\left(\frac{j}{\sum_{\ell\neq j}\mathbb{I}_{N_i^\ell>0}+1}\right)\right] \leq \mathbb{P}(N_i^j=0)\mathbb{E}\left[\log\left(\frac{j}{\sum_{\ell\neq j}\mathbb{I}_{N_i^\ell>0}+1}\right)\right] .$$

Resorting to Jensen's inequality and noticing that $\sum_{\ell\neq j}\mathbb{I}_{N_i^\ell>0}+1 \geq K_i$, we get

$$\mathbb{E}[C] \leq 2\mathbb{E}[M_{i,0}] + \sum_{j\geq 1} p_j \mathbb{P}(N_i^j=0)\log\mathbb{E}\left[\frac{j}{K_i}\right] .$$

Using Lemma 14, iii), we obtain

$$\mathbb{E}[C] \leq (2+\log 3)\,\mathbb{E}[M_{i,0}] + \sum_{j\geq 1} p_j \mathbb{P}(N_i^j=0)\log\left(\frac{j}{\mathbb{E}K_i}\right) . \tag{9}$$

Note that Proposition 12 follows from (8) and (9), combined with Lemma 14, i.

The next step consists in assuming that the source belongs to an envelope class $\mathcal{C}(f)$.

Lemma 17. *Let f be an envelope function and assume $P\in\mathcal{C}(f)$. Then*

i)

$$\sum_{j\geq 1} p_j\mathbb{P}(N_i^j=0)\log\left(\frac{j}{\mathbb{E}K_i}\right) \leq 2\frac{\mathbb{E}_f K_i+\ell_f}{i} + \sum_{j\geq \vec{\nu}_f(1/i)} f_j \log\left(\frac{j}{\vec{\nu}_f(1/i)}\right) .$$

ii)

$$\mathbb{E}[D] \leq 2\sum_{j\geq\vec{\nu}_f(1/i)} f_j \log\left(\log(j)+1\right) + 2\frac{\mathbb{E}_f K_i+\ell_f}{i}\log\left(\log\left(\vec{\nu}_f(1/i)\right)+1\right) .$$

Finally, in a last step, the regular assumption on the envelope allows us to evaluate each of the terms involved in the instantaneous redundancy.

Lemma 18. *Let $P\in\mathcal{C}(f)$ and assume that $\vec{\nu}_f(1/\cdot)\in \mathrm{RV}_\alpha$, for $\alpha\in[0,1[$. For all $\varepsilon>0$, there exists $i_0\in\mathbb{N}$ such that, for all $i\geq i_0$*

i)

$$\sum_{j\geq\vec{\nu}_f(1/i)} f_j\log\frac{j}{\vec{\nu}_f(1/i)} \leq \frac{3-\alpha}{(1-\alpha)^2}\,\frac{\vec{\nu}_f(1/i)}{i} .$$

ii)

$$\mathbb{E}[D] \leq 2(1+\varepsilon)\left(\Gamma(1-\alpha)+\frac{1}{1-\alpha}\right)\frac{\log\log(i)\vec{\nu}_f(1/i)}{i} .$$

As the bound on $\mathbb{E}[D]$ in Lemma 18 is much larger than $\vec{\nu}_f(1/i)/i$, we finally obtain that there exists $i_0\in\mathbb{N}$ that may depend on the envelope f such that for all $i\geq i_0$,

$$\mathbb{E}_P\left[\log\frac{P(X_{i+1})}{Q(X_{i+1}|X_{1:i})}\right] \leq c_\alpha\frac{\log\log(i)\vec{\nu}_f(1/i)}{i} ,$$

with $c_\alpha = 3\left(\Gamma(1-\alpha)+\frac{1}{1-\alpha}\right)$.

Hence, for $n\geq i_0$, we obtain

$$\overline{R}(Q_n,\mathcal{C}(f)^n) \leq C(i_0,f) + c_\alpha\log\log(n)\sum_{i=1}^n\frac{\vec{\nu}_f(1/i)}{i} ,$$

which establishes the upper bound in Theorem 6.

■

VII. Proofs

A. Proof of Theorem 4

i) As announced, the upper bound follows from bounds of [1] on the minimax regret defined as

$$\widehat{R}(\mathcal{C}(f)^n) = \inf_{Q_n} \sup_{P\in\mathcal{C}(f)} \sup_{x_{1:n}\in\mathcal{X}^n} \log \frac{P^n(x_{1:n})}{Q_n(x_{1:n})}\,.$$

Clearly, $\overline{R}\left(\mathcal{C}(f)^n\right) \le \widehat{R}\left(\mathcal{C}(f)^n\right)$. Now Theorem 14 of [1] states that

$$\widehat{R}\left(\mathcal{C}(f)^n\right) \le 1 + \widehat{R}\left(\mathcal{C}(f)^{\mathcal{P}(n)}\right) \;=\; 1 + \widehat{R}\left(\tau(\mathcal{C}(f)^{\mathcal{P}(n)})\right)\,,$$

Using (6) and the fact that for $j < \ell_f$, we still have the crude bound $p_j \le 1$,

$$\tau(\mathcal{C}(f)^{\mathcal{P}(n)}) \subset \left(\prod_{j=1}^{\ell_f-1} \mathcal{P}^\star(n)\right) \prod_{j=\ell_f}^{\infty} \mathcal{P}^\star(nf_j)\,.$$

As enlarging the class can not reduce the regret, and as, for two classes $\mathcal{C}_1, \mathcal{C}_2$, $\widehat{R}(\mathcal{C}_1\times\mathcal{C}_2) = \widehat{R}(\mathcal{C}_1) + \widehat{R}(\mathcal{C}_2)$, we have

$$\widehat{R}\left(\tau(\mathcal{C}(f)^{\mathcal{P}(n)})\right) \le \ell_f \widehat{R}\left(\mathcal{P}^\star(n)\right) + \sum_{j\ge\ell_f} \widehat{R}\left(\mathcal{P}^\star(nf_j)\right)\,.$$

Lemma 17 of [1] provide us with the following bounds: if $\lambda \le 1$,

$$\widehat{R}\left(\mathcal{P}^\star(\lambda)\right) = \log\left(2 - \mathrm{e}^{-\lambda}\right) \le \log(\mathrm{e})\lambda\,,$$

and, if $\lambda > 1$,

$$\widehat{R}\left(\mathcal{P}^\star(\lambda)\right) \le \log\left(\sqrt{\frac{2\lambda}{\pi}} + 2\right)\,.$$

For aesthetic purposes (getting the common $\log(\mathrm{e})$ constant in front of each terms), we find it convenient to notice that the bound above can be very slightly improved to

$$\widehat{R}\left(\mathcal{P}^\star(\lambda)\right) \le \log\left(\sqrt{\frac{2\lambda}{\pi}} + \frac{3}{2}\right)\,.$$

As $\frac{3}{2} + \sqrt{\frac{2}{\pi}} \le \mathrm{e}$,

$$\widehat{R}\left(\mathcal{C}(f)^n\right) \le O(\ell_f \log n) + \sum_{j, f_j\ge 1/n} \log\left(\mathrm{e}\sqrt{nf_j}\right) + \log(\mathrm{e}) \sum_{j, f_j\le 1/n} nf_j\,.$$

Now, using the integral representation and integrating by parts gives

$$\begin{aligned}\sum_{f_j\ge 1/n} \log\left(\mathrm{e}\sqrt{nf_j}\right) &= \log(\mathrm{e})\vec{\nu}_f(1/n) + \int_{1/n}^1 \frac{\log(nx)}{2}\nu_f(\mathrm{d}x)\\ &= \log(\mathrm{e})\left(\vec{\nu}_f(1/n) + \int_{1/n}^1 \frac{\vec{\nu}_f(x)}{2x}\mathrm{d}x\right)\,.\end{aligned}$$

Also, we may write $\sum_{j,f_j\le 1/n} f_j = \nu_{1,f}[0,1/n]$, which gives the desired upper bound.

ii) Let $m = n - n^{2/3}$. Thanks to Proposition 9, for n large enough

$$\overline{R}\left(\mathcal{C}(f)^n\right) \ge \overline{R}\left(\mathcal{C}(f)^{\mathcal{P}(m)}\right) - 1\,.$$

Now, by Lemmas 10 and 11,

$$\overline{R}\left(\mathcal{C}(f)^{\mathcal{P}(m)}\right) \ge \sum_{j=\ell_f}^{\infty} \overline{R}\left(\mathcal{P}^\star(mf_j)\right) \ge \sum_{j,\, f_j\ge 1/m} \overline{R}\left(\mathcal{P}^\star(mf_j)\right) \ge \sum_{j,\, f_j\ge 1/m} \left(\frac{\log(mf_j)}{2} - 5\right)\,.$$

Using the integral representation and integrating by parts,

$$\begin{aligned}\sum_{j,\, f_j \geq 1/m} \frac{\log(mf_j)}{2} &= \frac{\log(\mathrm{e})}{2} \int_{1/m}^{1} \ln(mx) \nu_f(\mathrm{d}x) \\ &= \frac{\log(\mathrm{e})}{2} \left([-\vec{\nu}_f(x) \ln(mx)]_{1/m}^{1} + \int_{1/m}^{1} \frac{\vec{\nu}_f(x)}{x} \mathrm{d}x \right) \\ &= \log(\mathrm{e}) \int_{1}^{m} \frac{\vec{\nu}_f(1/t)}{2t} \mathrm{d}t \,.\end{aligned}$$

We obtain

$$\overline{R}\left(\mathcal{C}(f)^n\right) \geq \log(\mathrm{e}) \int_{1}^{m} \frac{\vec{\nu}_f(1/t)}{2t} \mathrm{d}t - 5\vec{\nu}_f(1/m) - 1 \,. \tag{10}$$

Note that the lower bound (10) might be irrelevant. For instance, when $\vec{\nu}_f(1/\cdot) \in \mathrm{RV}_\alpha$ with $\alpha > \log(\mathrm{e})/10$, the right-hand side becomes negative. However, an order-optimal lower bound in the case of heavy-tailed envelopes (that is, when $\alpha > 0$) was established in [10]: the expected redundancy of a class $\mathcal{C}(f)^n$ is lower-bounded by the expected number of distinct symbols in a sample of size n drawn according to the envelope distribution $(f_j)_{j\geq 1}$. More precisely, for all envelope function f, there exists a constant $c_f \geq 0$ such that, for all $n \geq 1$, $\overline{R}\left(\mathcal{C}(f)^n\right) \geq \mathbb{E}_f[K_n] - c_f$.

B. Proofs of Section IV

Proof of Proposition 8: We have

$$\overline{R}(\mathcal{C}^{\mathcal{P}(n)}) = \inf_{Q \in \mathcal{M}_1(\mathcal{X}^*)} \sup_{P \in \mathcal{C}} D(P^{\mathcal{P}(n)}, Q).$$

Let $P \in \mathcal{C}$ and $Q \in \mathcal{M}_1(\mathcal{X}^*)$, with $\mathcal{X}^* = \cup_{k\geq 1} \mathcal{X}^k$. The distribution Q can be written as $Q = \sum_{k\geq 1} q(k) Q_k$, where $(q(k))_{k\geq 1}$ is a probability distribution over $\mathbb{N}$, and, for each $k \geq 1$, Q_k is a distribution over $\mathcal{C}^k$. Hence

$$\begin{aligned}D(P^{\mathcal{P}(n)}, Q) &= \sum_{k\geq 1} \mathbb{P}(N = k) \sum_{x \in \mathcal{C}^k} P^k(x) \log \frac{\mathbb{P}(N=k) P^k(x)}{q(k) Q_k} \\ &= D(\mathcal{P}(n), q) + \sum_{k\geq 1} \mathbb{P}(N = k) D(P^k, Q_k) \,. \end{aligned} \tag{11}$$

Maximizing in P and minimizing in $(q(k))_{k\geq 1}$ and $(Q_k)_{k\geq 0}$, we get

$$\overline{R}(\mathcal{C}^{\mathcal{P}(n)}) = \inf_{(q(k))} D(\mathcal{P}(n), q) + \inf_{(Q_k)} \sup_{P\in\mathcal{C}} \sum_{k\geq 1} \mathbb{P}(N=k) D(P^k, Q_k) \,.$$

The first term is equal to zero for $q = \mathcal{P}(n)$, implying that the distribution Q which achieves the minimax redundancy is also a Poisson mixture. Hence

$$\overline{R}(\mathcal{C}^{\mathcal{P}(n)}) = \inf_{(Q_k)} \sup_{P\in\mathcal{C}} \sum_{k\geq 1} \mathbb{P}(N=k) D(P^k, Q_k) \,.$$

■

Proof of Proposition 9: We start with the lower bound on $\overline{R}(\mathcal{C}^n)$. Let $m = n - n^{2/3}$, and let M be a Poisson random variable with mean m. By Proposition 8,

$$\begin{aligned}\overline{R}(\mathcal{C}^{\mathcal{P}(m)}) &= \inf_{(Q_k)} \sup_{P\in\mathcal{C}} \sum_{k\geq 0} \mathbb{P}(M=k) D(P^k, Q_k) \\ &\leq \inf_{(Q_k)} \sum_{k\geq 0} \mathbb{P}(M=k) \sup_{P\in\mathcal{C}} D(P^k, Q_k)\end{aligned}$$

Using the fact that the sequence $(\overline{R}(\mathcal{C}^k))_{k\geq 0}$ is increasing and sub-additive (see Proposition 7), we have

$$\begin{aligned}\overline{R}(\mathcal{C}^{\mathcal{P}(m)}) &\leq \overline{R}(\mathcal{C}^n) + \sum_{k>n} \mathbb{P}(M=k) \left(\overline{R}(\mathcal{C}^k) - \overline{R}(\mathcal{C}^n)\right) \\ &\leq \overline{R}(\mathcal{C}^n) + \sum_{k>n} \mathbb{P}(M=k) \overline{R}(\mathcal{C}^{k-n}) \\ &\leq \overline{R}(\mathcal{C}^n) + \overline{R}(\mathcal{C}) \sum_{k>n} \mathbb{P}(M=k)(k-n) \,.\end{aligned}$$

Resorting to Lemma 27 in Appendix E, we have

$$\begin{aligned}\sum_{k>n}\mathbb{P}(M=k)(k-n) &= \mathbb{E}\left[(M-n)\mathbb{I}_{\{M>n\}}\right] = \int_0^\infty \mathbb{P}\left(M-m>t+n^{2/3}\right)\mathrm{d}t\\ &\leq \int_0^n \mathrm{e}^{-\frac{n^{4/3}}{2(m+n^{2/3})}}\mathrm{d}t + \int_n^\infty \mathrm{e}^{-\frac{t^2}{6t}}\mathrm{d}t\\ &\leq n\mathrm{e}^{-n^{1/3}/2} + 6\mathrm{e}^{-n/6} \underset{n\to\infty}{\longrightarrow} 0\,.\end{aligned}$$

This establishes the lower bound on $\overline{R}(\mathcal{C}^n)$ in Proposition 9.

Let us now establish the upper bound. Let now $m = n + n^{2/3}$ and M be a Poisson random variable with mean m. Using the Bayesian representation of the minimax redundancy (see Proposition 7), we have

$$\overline{R}(\mathcal{C}^{\mathcal{P}(m)}) = \sup_{\pi\in\mathcal{M}_1(\mathcal{C})}\inf_{Q\in\mathcal{M}_1(\mathcal{X}^*)}\int D(P^{\mathcal{P}(m)},Q)\mathrm{d}\pi(P)\,.$$

Fix $\pi\in\mathcal{M}_1(\mathcal{C})$. Resorting to Equation (11) in the proof of Proposition 8, we have

$$\begin{aligned}&\inf_{Q\in\mathcal{M}_1(\mathcal{X}^*)}\int D(P^{\mathcal{P}(m)},Q)\mathrm{d}\pi(P)\\ &\quad= \inf_{(Q_k),(q(k))}\int\left(D(\mathcal{P}(m),q)+\sum_{k\geq 0}\mathbb{P}(M=k)D(P^k,Q_k)\right)\mathrm{d}\pi(P)\\ &\quad= \inf_{(Q_k)}\int\sum_{k\geq 0}\mathbb{P}(M=k)D(P^k,Q_k)\mathrm{d}\pi(P)\\ &\quad= \sum_{k\geq 0}\mathbb{P}(M=k)\inf_{Q_k}\int D(P^k,Q_k)\mathrm{d}\pi(P)\,.\end{aligned}$$

We claim that the sequence $\left(\inf_{Q_k}\int D(P^k,Q_k)\mathrm{d}\pi(P)\right)_{k\geq 0}$ is increasing with respect to k. Indeed, let $k\geq 0$ and let $Q_{k+1}\in\mathcal{M}_1(\mathcal{X}^k)$. Denote by $Q_{k+1}^{(k)}$ its restriction to the first k symbols. Then, for all $P\in\mathcal{M}_1(\mathcal{X})$, $D(Q_{k+1},P^{k+1})\geq D(P^k,Q_{k+1}^{(k)})$. Hence for all Q_{k+1} there exist $Q_k\in\mathcal{M}_1(\mathcal{X}^k)$ such that $\int D(P^k,Q_k)\mathrm{d}\pi(P)\leq\int D(P^{k+1},Q_{k+1})\mathrm{d}\pi(P)$, which gives the desired result. We get

$$\begin{aligned}\overline{R}(\mathcal{C}^{\mathcal{P}(m)}) &\geq \sup_\pi\sum_{k\geq n}\mathbb{P}(M=k)\inf_{Q_k}\int D(P^k,Q_k)\mathrm{d}\pi(P)\\ &\geq \mathbb{P}(M\geq n)\overline{R}(\mathcal{C}^n)\,.\end{aligned}$$

Now, using again Lemma 27, we have

$$\mathbb{P}(M\geq n)\geq 1-\exp\left(-\frac{n^{4/3}}{2m}\right)\underset{n\to\infty}{\longrightarrow}1\,,$$

which concludes the proof. ∎

Proof of Lemma 11: Using the Bayesian representation of the minimax redundancy, we have

$$\overline{R}(\mathcal{P}^\star(\lambda)) = \sup_{\pi\in\mathcal{M}_1([0,\lambda])}\int D(\mathcal{P}(\mu),\mathcal{P}_\pi)\mathrm{d}\pi(\mu)\,,$$

where $\mathcal{P}_\pi = \int\mathcal{P}(\mu)\mathrm{d}\pi(\mu)$. In particular, taking π equal to the uniform distribution over $[0,\lambda]$, we get

$$\overline{R}(\mathcal{P}^\star(\lambda))\geq\int_0^\lambda\frac{1}{\lambda}\sum_{k\geq 0}\mathbb{P}(\mathcal{P}(\mu)=k)\log\frac{\mathbb{P}(\mathcal{P}(\mu)=k)}{\mathbb{P}(\mathcal{P}_\pi=k)}\mathrm{d}\mu\,.$$

We have

$$\mathbb{P}(\mathcal{P}_\pi=k)=\frac{1}{\lambda}\int_0^\lambda\frac{\mathrm{e}^{-\mu}\mu^k}{k!}\mathrm{d}\mu = \frac{\mathbb{P}(\mathcal{P}(\lambda)>k)}{\lambda}\leq\frac{1}{\lambda}\,.$$

Hence

$$\overline{R}(\mathcal{P}^\star(\lambda))\geq\log\lambda-\frac{1}{\lambda}\int_0^\lambda H(\mathcal{P}(\mu))\mathrm{d}\mu\,.$$

Using Stirling's bound $k! \leq \mathrm{e}^{1/12k} \left(\frac{k}{\mathrm{e}}\right)^k \sqrt{2\pi k}$, we have, for all $\mu \in [0, \lambda]$,

$$\begin{aligned}\frac{H(\mathcal{P}(\mu))}{\log(\mathrm{e})} &= \mu - \mu \ln \mu + \sum_{k\geq 0} \frac{\mathrm{e}^{-\mu}\mu^k}{k!} \ln(k!) \\ &\leq \mu - \mu \ln \mu + \sum_{k\geq 1} \frac{\mathrm{e}^{-\mu}\mu^k}{k!} \left(k \ln k - k + \frac{\ln(2\pi k)}{2} + \frac{1}{12k} \right) \\ &\leq \sum_{k\geq 1} \frac{\mathrm{e}^{-\mu}\mu^k}{k!} \left(k \ln k + \frac{\ln(2\pi k)}{2} \right) - \mu \ln \mu + \frac{1}{12} \,.\end{aligned}$$

We use Jensen's inequality to obtain

$$\sum_{k\geq 1} \frac{\mathrm{e}^{-\mu}\mu^k}{k!} k \ln k = \mu \sum_{k\geq 0} \frac{\mathrm{e}^{-\mu}\mu^k}{k!} \ln(k+1) \;\leq\; \mu \ln(1+\mu) \,,$$

and

$$\sum_{k\geq 1} \frac{\mathrm{e}^{-\mu}\mu^k}{k!} \ln k \leq (1 - \mathrm{e}^{-\mu}) \ln \left(\frac{\mu}{1 - \mathrm{e}^{-\mu}} \right) \;\leq\; \ln \mu + \frac{1}{\mathrm{e}} \,,$$

where the last inequality is due to the fact that the function $x \mapsto x \ln x$ is larger than $-1/\mathrm{e}$ for all $x \geq 0$. We get

$$\begin{aligned}\frac{H(\mathcal{P}(\mu))}{\log(\mathrm{e})} &\leq \frac{\ln \mu}{2} + \mu \ln \left(1 + \frac{1}{\mu} \right) + \frac{\ln(2\pi)}{2} + \frac{1}{2\mathrm{e}} + \frac{1}{12} \\ &\leq \frac{\ln \mu}{2} + 3 \,,\end{aligned}$$

which is smaller than 3 for $\mu \leq 1$. Hence

$$\begin{aligned}\overline{R}(\mathcal{P}^\star(\lambda)) &\geq \log \lambda - 3 \log(\mathrm{e}) - \frac{1}{\lambda} \int_1^\lambda \frac{\log \mu}{2} \mathrm{d}\mu \\ &\geq \frac{\log \lambda}{2} - 5 \,.\end{aligned}$$

■

C. Proofs of Section VI

Proof of Lemma 14:

i) Observe that

$$\mathbb{E}M_{n,0} = \int_0^1 x(1-x)^n \nu(\mathrm{d}x) \leq \frac{1}{n} \int_0^1 nx(1-x)^{n-1} \nu(\mathrm{d}x) = \frac{\mathbb{E}K_{n,1}}{n} \leq \frac{\mathbb{E}K_n}{n} \,.$$

The relation between $\mathbb{E}K_n$ and $\mathbb{E}_f K_n$ is easily obtained by noticing that the function $x \mapsto 1 - (1-x)^n$ is increasing on $[0, 1]$, which gives

$$\mathbb{E}K_n = \sum_{j\geq 1} (1 - (1-p_j)^n) \;\leq\; \ell_f + \sum_{j\geq 1} (1 - (1-f_j)^n) \;=\; \ell_f + \mathbb{E}_f K_n \,.$$

ii) The expected number of distinct symbols in a message of length i is never much less that the number of symbols with probability larger than $1/i$:

$$\mathbb{E}K_i \geq \sum_{j \leq \vec{\nu}(1/i)} \left(1 - \left(1 - p_j\right)^i\right) \geq \frac{\mathrm{e}-1}{\mathrm{e}} \vec{\nu}(1/i) \,.$$

The upper bound is obtained in a similar way

$$\begin{aligned}\mathbb{E}K_i &= \sum_{j\geq 1} \left(1 - (1-p_j)^i\right) \\ &\leq \vec{\nu}(1/i) + \sum_{j > \vec{\nu}(1/i)} \left(1 - (1-p_j)^i\right) \\ &\leq \vec{\nu}(1/i) + \sum_{j > \vec{\nu}(1/i)} i p_j \\ &= \vec{\nu}(1/i) + i \nu_1(0, 1/i)\end{aligned}$$

where the second inequality comes from $1-x \le (1-x/i)^i$ for $0 \le x \le 1$ and $i \ge 1$.

iii) The quantity $\mathbb{E}1/K_i$ satisfies a quadratic inequality:

$$\begin{aligned}\mathbb{E}\left[\frac{1}{K_i}\right] - \frac{1}{\mathbb{E}K_i} &= \mathbb{E}\left[\frac{\mathbb{E}[K_i]-K_i}{\sqrt{\mathbb{E}K_i}} \times \frac{1}{K_i\sqrt{\mathbb{E}K_i}}\right] \\ &\le \mathbb{E}\left[\frac{(\mathbb{E}[K_i]-K_i)^2}{\mathbb{E}K_i}\right]^{1/2} \times \mathbb{E}\left[\frac{1}{K_i^2\mathbb{E}K_i}\right]^{1/2} \\ &\le \frac{(\mathbb{E}K_{i,1})^{1/2}}{\mathbb{E}K_i} \times \mathbb{E}\left[\frac{1}{K_i}\right]^{1/2}\end{aligned}$$

where the first inequality follows by invoking Cauchy-Schwarz inequality, the first factor in the penultimate line is upper bounded by 1 thanks to the fact that $\mathrm{var}(K_i) \le \mathbb{E}K_{i,1}$ [4], while the second factor is crudely upper bounded using the fact that $K_i \ge 1$ entails $1/K_i^2 \le 1/K_i$. Solving the inequality leads to

$$\begin{aligned}\mathbb{E}\left[\frac{1}{K_i}\right] &\le \frac{1}{\mathbb{E}K_i}\left(1 + \frac{\mathbb{E}K_{i,1}}{2\mathbb{E}K_i} + \left(\frac{\mathbb{E}K_{i,1}}{\mathbb{E}K_i}\right)^{1/2}\sqrt{1+\frac{\mathbb{E}K_{i,1}}{4\mathbb{E}K_i}}\right) \\ &\le \frac{1}{\mathbb{E}K_i}\left(1 + \frac{\mathbb{E}K_{i,1}}{\mathbb{E}K_i} + \left(\frac{\mathbb{E}K_{i,1}}{\mathbb{E}K_i}\right)^{1/2}\right) \\ &\le \frac{1}{\mathbb{E}K_i}\left(1 + 2\left(\frac{\mathbb{E}K_{i,1}}{\mathbb{E}K_i}\right)^{1/2}\right) \\ &\le \frac{3}{\mathbb{E}K_i}\,.\end{aligned}$$

■

Proof of Lemma 17:

i) We have

$$\begin{aligned}\sum_{j\ge1} p_j\mathbb{P}(N_i^j=0)\log\left(\frac{j}{\mathbb{E}[K_i]}\right) &= \sum_{j\ge1} p_j\mathbb{P}(N_i^j=0)\log\left(\frac{j}{\mathbb{E}_f[K_i]}\right) + \mathbb{E}[M_{i,0}]\log\frac{\mathbb{E}_f[K_i]}{\mathbb{E}[K_i]} \\ &\le \sum_{j\ge1} p_j\mathbb{P}(N_i^j=0)\log\left(\frac{j}{\mathbb{E}_f[K_i]}\right) + \frac{\mathbb{E}K_i}{i}\log\frac{\mathbb{E}_f[K_i]}{\mathbb{E}[K_i]} \\ &\le \sum_{j\ge1} p_j\mathbb{P}(N_i^j=0)\log\left(\frac{j}{\mathbb{E}_f[K_i]}\right) + \frac{\mathbb{E}_f[K_i]}{2i}\,,\end{aligned}$$

where the first inequality follows from $\mathbb{E}[M_{i,0}] \le \mathbb{E}[K_i]/i$ (see Lemma 14), and the second inequality follows from the fact that $x \mapsto x\log\frac{y}{x}$ achieves its maximum $y/2$ at $x = y/2$.
By Lemma 14, ii, we have $\mathbb{E}_f K_i \ge \frac{\mathrm{e}-1}{\mathrm{e}}\vec{\nu}_f(1/i)$. Hence

$$\sum_{j\ge1} p_j\mathbb{P}(N_i^j=0)\log\left(\frac{j}{\mathbb{E}_f[K_i]}\right) \le \log\frac{\mathrm{e}}{\mathrm{e}-1}\,\mathbb{E}[M_{i,0}] + \sum_{j\ge1} p_j\log\left(\frac{j}{\vec{\nu}_f(1/i)}\right)\,.$$

Now, for $j < \vec{\nu}_f(1/i)$, the summands are negative and we simply omit them to get

$$\sum_{j\ge1} p_j\log\left(\frac{j}{\vec{\nu}_f(1/i)}\right) \le \sum_{j\ge\vec{\nu}_f(1/i)} f_j\log\left(\frac{j}{\vec{\nu}_f(1/i)}\right)\,,$$

where we used $p_j \le f_j$, provided that $j \ge \ell_f$.

ii) Decomposing the sum according to $\vec{\nu}_f(1/i)$ and using the fact that $\mathbb{E}[M_{i,0}] \le (\mathbb{E}_f K_i + \ell_f)/i$ readily yield the desired bound on $\mathbb{E}[D]$.

■

Proof of Lemma 18:

i) Using the integral representation, we have

$$\sum_{j\ge\vec{\nu}_f(1/i)} f_j\log\left(\frac{j}{\vec{\nu}_f(1/i)}\right) = \int_0^{1/i} x\log\frac{\vec{\nu}_f(x)}{\vec{\nu}_f(1/i)}\nu_f(\mathrm{d}x)\,.$$

By the regular variation assumption on $\vec{\nu}_f(1/\cdot)$, the Potter-Drees Inequality (see section A) implies that: for all $\varepsilon, \delta > 0$,

$\exists i_0 \in \mathbb{N}$ such that, for all $i \geq i_0$, for all $x \in (0, 1/i]$,

$$\frac{\vec{\nu}_f(x)}{\vec{\nu}_f(1/i)} \leq \left(\frac{1}{xi}\right)^{\alpha} + \varepsilon \left(\frac{1}{xi}\right)^{\alpha+\delta} .$$

Taking crudely $\varepsilon = 1$, $\delta = 1 - \alpha$ and bounding α by 1, we obtain that for i large enough and $x \in (0, 1/i]$,

$$\frac{\vec{\nu}_f(x)}{\vec{\nu}_f(1/i)} \leq \frac{2}{xi} .$$

Hence

$$\int_0^{1/i} x \log \frac{\vec{\nu}_f(x)}{\vec{\nu}_f(1/i)} \nu_f(\mathrm{d}x) \leq \nu_{1,f}[0, 1/i] + \log(\mathrm{e}) \int_0^{1/i} x \ln \left(\frac{1}{xi}\right) \nu_f(\mathrm{d}x) .$$

By Fubini's Theorem,

$$\begin{aligned} \int_0^{1/i} x \ln \left(\frac{1}{xi}\right) \nu_f(\mathrm{d}x) &= \int_0^{1/i} x \int_1^{\frac{1}{xi}} \frac{1}{t} \mathrm{d}t \, \nu_f(\mathrm{d}x) \\ &= \int_1^{\infty} \frac{1}{t} \int_0^{\frac{1}{ti}} x \nu_f(\mathrm{d}x) \mathrm{d}t \; = \; \int_i^{\infty} \frac{\nu_{1,f}[0, 1/t]}{t} \mathrm{d}t . \end{aligned}$$

Now, Lemma 16 implies that, for all $\varepsilon > 0$, for i large enough and for all $t \geq i$, $\nu_{1,f}[0, 1/t] \leq \frac{(\alpha+\varepsilon)\vec{\nu}_f(1/t)}{(1-\alpha)t}$, and by Karamata's Theorem (see section A),

$$\int_i^{\infty} \frac{\vec{\nu}_f(1/t)}{t^2} \mathrm{d}t \sim \frac{\vec{\nu}_f(1/i)}{(1-\alpha)i} .$$

Hence, for i large enough,

$$\sum_{j \geq \vec{\nu}_f(1/i)} f_j \log \left(\frac{j}{\vec{\nu}_f(1/i)}\right) \leq \frac{3-\alpha}{(1-\alpha)^2} \times \frac{\vec{\nu}_f(1/i)}{i} .$$

ii) By Lemma 17,

$$\mathbb{E}[D] \leq 2 \sum_{j \geq \vec{\nu}_f(1/i)} f_j \log \left(\log(j) + 1\right) + 2 \frac{\mathbb{E}_f K_i + \ell_f}{i} \log \left(\log \left(\vec{\nu}_f(1/i)\right) + 1\right) . \tag{12}$$

As $\vec{\nu}_f(x) \ll 1/x$ when $x \to 0$ (which is true as soon as the support is infinite, see [24]), for i large enough, we have $\log(\vec{\nu}_f(x)) + 1 \leq \log(1/x)$ for all $x \in]0, 1/i]$. Hence, the first term in the right-hand side of (12) can be controlled as follows,

$$\begin{aligned} \sum_{j \geq \vec{\nu}_f(1/i)} f_j \log \left(\log(j) + 1\right) &= \int_0^{1/i} x \log \left(\log(\vec{\nu}_f(x)) + 1\right) \nu_f(\mathrm{d}x) \\ &\leq \int_0^{1/i} x \log \log \left(\frac{1}{x}\right) \nu_f(\mathrm{d}x) \\ &= \left[-\vec{\nu}_f(x) x \log \log \left(\frac{1}{x}\right)\right]_0^{1/i} + \int_0^{1/i} \left(\log \log \left(\frac{1}{x}\right) + \frac{\log(\mathrm{e})}{\ln x}\right) \vec{\nu}_f(x) \mathrm{d}x \\ &\leq \int_0^{1/i} \log \log \left(\frac{1}{x}\right) \vec{\nu}_f(x) \mathrm{d}x \; = \; \int_i^{\infty} \frac{\log \log(t) \vec{\nu}_f(1/t)}{t^2} \mathrm{d}t . \end{aligned}$$

We used that, as $\alpha < 1$, the limit of $\vec{\nu}_f(x) x \log \log(1/x)$ as $x \to 0$ is equal to 0. Now, the function $t \mapsto \frac{\log \log(t) \vec{\nu}_f(1/t)}{t^2}$ belongs to $\mathrm{RV}_{\alpha-2}$. Hence, by Karamata's Theorem (see section A),

$$\int_i^{\infty} \frac{\log \log(t) \vec{\nu}_f(1/t)}{t^2} \mathrm{d}t \underset{i \to \infty}{\sim} \frac{\vec{\nu}_f(1/i)}{(1-\alpha)i} \log \log(i) .$$

As for the second term in the right-hand side of (12), we use $\vec{\nu}_f(1/i) \ll i$ and $\mathbb{E}_f K_i \sim \Gamma(1-\alpha)\vec{\nu}_f(1/i)$ and obtain that, for all $\varepsilon > 0$, there exists $i_0 \in \mathbb{N}$ such that, for all $i \geq i_0$

$$\mathbb{E}[D] \leq 2(1+\varepsilon) \left(\Gamma(1-\alpha) + \frac{1}{1-\alpha}\right) \frac{\log \log(i) \vec{\nu}_f(1/i)}{i} .$$

■

REFERENCES

[1] J. Acharya, A. Jafarpour, A. Orlitsky, and A.T. Suresh. Universal compression of envelope classes: Tight characterization via Poisson sampling. *ArXiv e-prints*, abs/1405.7460, 2014.

[2] C.W. Anderson. Extreme value theory for a class of discrete distributions with applications to some stochastic processes. *J. Appl. Probability*, 7:99–113, 1970.

[3] S. Arlot. Model selection by resampling penalization. *Electron. J. Stat.*, 3:557–624, 2009.

[4] A. Ben-Hamou, S. Boucheron, and M.I. Ohannessian. Concentration inequalities in the infinite urn scheme for occupancy counts and the missing mass, with applications. *Bernoulli*, 2017. to appear, ArXiv e-prints arXiv:1412.8652.

[5] N.H. Bingham, C.M. Goldie, and J.L. Teugels. *Regular variation*, volume 27. Cambridge University Press, 1989.

[6] D. Bontemps. Universal coding on infinite alphabets: exponentially decreasing envelopes. *IEEE Trans. Inform. Theory*, 57:1466–1478, 2011.

[7] D. Bontemps, S. Boucheron, and E. Gassiat. About adaptive coding on countable alphabets. *IEEE Trans. Inform. Theory*, 60:808–821, 2014.

[8] S. Boucheron, A. Garivier, and E. Gassiat. Coding over Infinite Alphabets. *IEEE Trans. Inform. Theory*, 55:358–373, 2009.

[9] S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford Univ. Press, 2013.

[10] S. Boucheron, E. Gassiat, and M.I. Ohannessian. About adaptive coding on countable alphabets: max-stable envelope classes. *IEEE Trans. Inform. Theory*, 61:4948–4967, 2015.

[11] O. Catoni. *Statistical learning theory and stochastic optimization*, volume 1851 of *Lecture Notes in Mathematics*. Springer-Verlag, 2004. Ecole d'Eté de Probabilités de Saint-Flour XXXI.

[12] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge Univ. Press, 2006.

[13] B.S. Clarke and A.R. Barron. Information-theoretic asymptotics of Bayes methods. *IEEE Trans. Inform. Theory*, 36: 453–471, 1990.

[14] B.S. Clarke and A.R Barron. Jeffrey's prior is asymptotically least favorable under entropy risk. *J. Stat. Planning and Inference*, 41:37–60, 1994.

[15] T. Cover and J. Thomas. *Elements of information theory*. John Wiley & sons, 1991.

[16] I. Csiszár and J. Körner. *Information Theory: Coding Theorems for Discrete Memoryless Channels*. Academic Press, 1981.

[17] L. de Haan and A. Ferreira. *Extreme value theory*. Springer-Verlag, 2006.

[18] D. Dubhashi and D. Ranjan. Balls and bins: A study in negative dependence. *Random Struct. & Algorithms*, 13:99–124, 1998.

[19] P. Elias. Universal codeword sets and representations of the integers. *IEEE Trans. Inform. Theory*, 21:194–203, 1975.

[20] M. Falahatgar, A. Jafarpour, A. Orlitsky, V. Pichapati, and A. Theertha Suresh. Universal Compression of Power-Law Distributions. *ArXiv e-prints*, 2015.

[21] A. Garivier. Redundancy of the context-tree weighting method on renewal and Markov renewal processes. *IEEE Trans. Inform. Theory*, 52:5579–5586, 2006.

[22] A. Garivier. A lower bound for the maximin redundancy in pattern coding. *Entropy*, 11:634–642, 2009.

[23] E. Gassiat. *Codage universel et identification d'ordre par sélection de modèles*. Société Mathématique de France, 2014.

[24] A. Gnedin, B. Hansen, and J. Pitman. Notes on the occupancy problem with infinitely many boxes: general asymptotics and power laws. *Probab. Surv.*, 4:146–171, 2007. doi: 10.1214/07-PS092.

[25] L. Györfi, M. Kohler, A. Krzyzak, and H. Walk. *A distribution-free theory of nonparametric regression*. Springer-Verlag, 2006.

[26] D. Haussler and M. Opper. Mutual information, metric entropy and cumulative relative entropy risk. *Annals of Statistics*, 25:2451–2492, 1997.

[27] P. Jacquet and W. Szpankowski. Analytical depoissonization and its applications. *Theoretical Computer Science*, 201: 1–62, 1998.

[28] J. Jiao, K. Venkat, Y. Han, and T. Weissman. Minimax estimation of functionals of discrete distributions. *IEEE Trans. Inform. Theory*, 61:2835–2885, 2015.

[29] S. Karlin. Central limit theorems for certain infinite urn schemes. *J. Math. Mech.*, 17:373–401, 1967.

[30] J.C. Kieffer. A unified approach to weak universal source coding. *IEEE Trans. Inform. Theory*, 24:674–682, 1978.

[31] N. Merhav and M. Feder. Universal prediction. *IEEE Trans. Inform. Theory*, 44:2124–2147, 1998.

[32] A. Orlitsky and N.P. Santhanam. Speaking of infinity. *IEEE Trans. Inform. Theory*, 50:2215–2230, 2004.

[33] A. Orlitsky, N.P. Santhanam, and J. Zhang. Always Good-Turing: asymptotically optimal probability estimation. *Science*, 302:427–431, 2003.

[34] A. Orlitsky, N.P. Santhanam, K. Viswanathan, and J. Zhang. On modeling profiles instead of values. In *Proc. of the 20th conf. on Uncertainty in Artificial Intelligence*, pages 426–435. AUAI Press, 2004.

[35] A. Orlitsky, N.P. Santhanam, and J. Zhang. Universal compression of memoryless sources over unknown alphabets. *IEEE Trans. Inform. Theory*, 50:1469–1481, 2004.
[36] J. Rissanen and Jr. G. G. Langdon. Arithmetic coding. *IBM J. Res. Develop.*, 23:149–162, 1979.
[37] B.Y. Ryabko. Twice-universal coding. *Problemy Peredachi Informatsii*, 20:24–28, 1984.
[38] N. Santhanam, V. Anantharam, A. Kavcic, and W. Szpankowski. Data-driven weak universal redundancy. In *IEEE Intern. Symp. Inform. Theory*, pages 1877–1881. IEEE, 2014.
[39] Gil I. Shamir. Universal Lossless Compression with Unknown Alphabets - The Average Case, 2006. URL `http://www.citebase.org/abstract?id=oai:arXiv.org:cs/0603068`.
[40] O. Shayevitz, R. Zamir, and M. Feder. Bounded expected delay in arithmetic coding. In *IEEE Intern. Symp. Inform. Theory*, pages 2662–2666, 2006.
[41] A.T. Suresh. *Statistical Inference over Large Domains*. PhD Dissertation, UCSD, 2016.
[42] W. Szpankowski and M.J. Weinberger. Minimax pointwise redundancy for memoryless models over large alphabets. *IEEE Trans. Inform. Theory*, 58:4094–4104, 2012.
[43] R.-E. Tarjan. *Data structures and network algorithms*, volume 44 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. (SIAM), 1983.
[44] A.B. Tsybakov. *Introduction à l'estimation non-paramétrique*, volume 41 of *Mathématiques & Applications*. Springer-Verlag, Berlin, 2004.
[45] G. Valiant and P. Valiant. Estimating the unseen: an n/log(n)-sample estimator for entropy and support size, shown optimal via new CLTs. In *Proc. of 43rd ACM Symp. on Theory Of Computing*, pages 685–694. ACM, 2011.
[46] F. M.J. Willems. The context-tree weighting method: extensions. *IEEE Trans. Inform. Theory*, 44:792–798, 1998.
[47] Q. Xie and A. R. Barron. Minimax redundancy for the class of memoryless sources. *IEEE Trans. Inform. Theory*, 43: 646–656, 1997.
[48] Q. Xie and A. R. Barron. Asymptotic minimax regret for data compression, gambling and prediction. *IEEE Trans. Inform. Theory*, 46:431–445, 2000.
[49] X. Yang and A. Barron. Large Alphabet Coding and Prediction through Poissonization and Tilting. In *The Sixth Workshop on Information Theoretic Methods in Science and Engineering*, Tokyo, August 2013.

APPENDIX A
REGULAR VARIATION

The determination of bounds on minimax redundancies and the analysis of the performance of the PC code for classes with regularly varying envelopes rely at some point on classical results from regular variation theory. For the sake of self-reference, these results are recalled here [See 5, 17, for proofs and refinements].

Theorem 19. (KARAMATA'S INTEGRATION THEOREM) *Let $f \in \mathrm{RV}_\alpha$ and assume that there exists $t_0 > 0$ such that f is positive and locally bounded on $[t_0, +\infty[$.*

(i) If $\alpha > -1$, then

$$\int_{t_0}^{t} f(s)\mathrm{d}s \underset{t\to+\infty}{\sim} \frac{tf(t)}{\alpha+1}\,.$$

(ii) If $\alpha < -1$, or $\alpha = -1$ and $\int_0^\infty f(s)\mathrm{d}s < \infty$, then

$$\frac{tf(t)}{\int_t^{+\infty} f(s)\mathrm{d}s} \underset{t\to+\infty}{\sim} -\alpha - 1\,.$$

Theorem 20. (POTTER-DREES INEQUALITIES.) *If $f \in \mathrm{RV}_\alpha$, then for all $\delta > 0$, there exists $t_0 = t_0(\alpha)$, such that for all $t, x\colon\ \min(t, tx) > t_0$,*

$$(1-\delta)x^\alpha \min\left(x^\delta, x^{-\delta}\right) \le \frac{f(tx)}{f(t)} \le (1+\delta)x^\alpha \max\left(x^\delta, x^{-\delta}\right).$$

In order to relate Karlin's infinite urn scheme setting and the setting of [10] in Appendix C, we need to compute asymptotic inverses of regularly varying functions. This is done using the notion of De Bruijn conjugacy.

Theorem 21 (DE BRUIJN CONJUGACY)**.** *[5, Proposition 1.5.15] Let $\ell \in \mathrm{RV}(0)$, then there exists a function $\ell^\# \in \mathrm{RV}(0)$ such that $\ell^\#(x)\ell(x\ell^\#(x)) \to 1$ and $\ell(x)\ell^\#(x\ell(x)) \to 1$ as $x \to \infty$. Any function satisfying these two relations is asymptotically equivalent to $\ell^\#$. The functions $(\ell, \ell^\#)$ are said to form a pair of De Bruijn conjugates.*

If $f = t^\gamma \ell(t)$ where ℓ is slowly varying, then $\ell^{1/\gamma}$ is also slowly varying and the function $y \mapsto y^{1/\gamma}\left(\ell^{1/\gamma}\right)^\#\left(y^{1/\gamma}\right)$ is an asymptotic inverse of f. This entails that for $\alpha > 0$, $f \in \mathrm{RV}_\alpha$, any asymptotic inverse of f is regularly varying with regular variation index $1/\alpha$.

While exploring source classes defined by slowly varying counting functions, a refined notion of regular variation due to de Haan may prove relevant.

Definition 22. A slowly varying function ℓ has the extended regular variation property if there exists a slowly varying function ℓ_0 such that for all $x > 0$

$$\lim_{t\to\infty} \frac{\ell(tx) - \ell(t)}{\ell_0(t)} = \log(x)\,,$$

this is summarized by $\ell \in \Pi_{\ell_0}$.

Slowly varying functions satisfying the extended regular variation property form a proper subset of RV_0. For example, $g(t) = \lfloor \log(t) \rfloor$ and $g(t) = \lfloor \log(t) \rfloor^2$ are slowly varying but do not satisfy the extended regular variation property, while $g(t) = \lfloor \log(t)^2 \rfloor \in \Pi_{2\log}$.

Appendix B
Infinite urn schemes and regular variation

The next theorem connects $\mathbb{E}K_n$ and the regular variation properties of the envelope (characterized by the regular variation of $\vec{\nu}$). For much more on the asymptotic behavior of occupancy counts, see Karlin [29], Gnedin, Hansen, and Pitman [24].

Theorem 23 ([29, 24])**.** *Assume that there exists a slowly varying function ℓ such that $\vec{\nu}(1/n) = n^\alpha \ell(n)$, with $\alpha \in [0,1]$ and ℓ slowly varying at infinity.*

i) If $0 < \alpha < 1$,

$$\mathbb{E}K_n \sim \Gamma(1-\alpha)\vec{\nu}(1/n)\,, \quad \textit{and} \quad \nu_1[0,1/n] \sim \frac{\alpha\vec{\nu}(1/n)}{(1-\alpha)n}\,.$$

ii) If $\alpha = 1$,

$$\mathbb{E}K_n \sim \mathbb{E}K_{n,1} \sim n\nu_1[0,1/n] \sim n\ell_1(n) \quad \textit{where} \quad \ell_1(n) = \int_n^\infty \frac{\ell(t)}{t}\mathrm{d}t \gg \ell(n)\,.$$

iii) If $\alpha = 0$,

$$\mathbb{E}K_n \sim \ell(n)\,, \quad \textit{and} \quad \nu_1[0,1/n] \ll \frac{\ell(n)}{n}\,.$$

If furthermore, $\vec{\nu}(1/\cdot) \in \Pi_{\ell_0}$, then $\nu_1[0,x] \sim x\ell_0(1/x)$ and $\ell(x) \sim \int_1^x u^{-1}\ell_0(u)\mathrm{d}u$.

Appendix C
Connections between counting function and tail quantile function

In [7, 10], censoring methods and minimax redundancy rates for envelope classes are described using the envelope tail quantile function (or a smothed version of it) $U(t) = \inf\{x : \overline{F}(x) \le \frac{1}{t}\}$ (where $\overline{F}$ denotes the envelope survival function). In order to clarify the connection between the performance of the PC code and the performance of the codes described in these two references, we relate the regular variation properties of the counting function $\vec{\nu}$ and the regular variation properties of the tail quantile function U.

For the sake of simplicity we assume that the probability mass function $(p_j)_{j\ge1}$ is decreasing and that no p_j is null. This assumption entails $\vec{\nu}(p_j) = j$ for $j \ge 1$.

For $t > 1$,

$$U(t) = \inf\left\{j : \textstyle\sum_{k>j} p_k \le \frac{1}{t}\right\} = \inf\left\{j : \nu_1(0,p_j) \le \frac{1}{t}\right\}.$$

This entails that U may be defined from $\vec{\nu}$ and ν_1,

$$U(t) = 1 + \vec{\nu}\left(\sup\left\{x : \nu_1(0,x) \le \tfrac{1}{t}\right\}\right)\,.$$

Note that $1/\sup\left\{x : \nu_1(0,x) \le \frac{1}{t}\right\} = \inf\left\{y : 1/\nu_1(0,1/y) \ge t\right\}$, so letting $h(y)$ be a shorthand for $1/\nu_1(0,1/y)$, $U(t) = 1 + \vec{\nu}(1/h^{\leftarrow}(t))$.

Lemma 24. *Assume that there exists a slowly varying function ℓ such that $\vec{\nu}(x) \sim x^{-\alpha}\ell(1/x)$, with $\alpha \in [0,1)$ and that the probability mass function $(p_j)_{j\ge1}$ is decreasing and that no p_j is null. Let $\gamma = \alpha/(1-\alpha)$. Then,*

i) U is asymptotically equivalent with a regularly varying function with index γ.

ii) If $\alpha \in (0,1)$, let $\widetilde{\ell}$ be a shorthand for $(1/\ell^{1+\gamma})^{\#}$. Then

$$U(t) \underset{t\to\infty}{\sim} \gamma^\gamma t^\gamma \widetilde{\ell}\left(t^{1+\gamma}\right)\,.$$

iii) If $\alpha = 0$ and $\ell \in \Pi_{\ell_0}$, let $\widetilde{\ell}_0$ be a shorthand for $(1/\ell_0)^{\#}$. Then U has the extended regular variation property, $U \in \Pi_{\widetilde{\ell}_0}$ and

$$U(t) \underset{t\to\infty}{\sim} \ell\left(t\widetilde{\ell}_0(t)\right)\,.$$

iv) If $\alpha = 0$ and the distribution is discrete log-concave, then $U(t) \underset{t\to\infty}{\sim} \vec{\nu}(1/t)$.

Proof:

If $\vec{\nu}(1/\cdot)$ has either positive regular variation index or has both the slow variation and the extended regular variation property, then $h^{\leftarrow}$ is the generalized inverse of the regularly varying function $1/\nu_1(0,1/\cdot)$ (See Theorem 23 above). As such, $h^{\leftarrow}$ is regularly varying, and so is U.

If $h(y) = 1/\nu_1(0,1/\cdot) \in \text{RV}_{1-\alpha}$ for $\alpha \in (0,1)$, then its generalized inverse is regularly varying with index $1+\gamma = 1/(1-\alpha)$ and, from Propositions 1.5.14 and 1.5.15 from [5],

$$h^{\leftarrow}(t) \underset{t\to\infty}{\sim} \gamma^{1+\gamma} t^{1+\gamma} \widetilde{\ell}\left(t^{1+\gamma}\right) .$$

and

$$\begin{aligned} U(t) &\underset{t\to\infty}{\sim} \gamma^{\gamma} t^{\gamma} \widetilde{\ell}\left(t^{1+\gamma}\right)^{\gamma/(1+\gamma)} \ell\left(t^{1+\gamma}\widetilde{\ell}\left(t^{1+\gamma}\right)\right) \\ &\underset{t\to\infty}{\sim} \gamma^{\gamma} t^{\gamma} \widetilde{\ell}\left(t^{1+\gamma}\right) . \end{aligned}$$

If $\vec{\nu}(1/\cdot) \in \Pi_{\ell_0}$ where ℓ_0 is slowly varying, then by Theorem 23 and conjugacy arguments

$$U(t) \sim \ell\left(t\widetilde{\ell}_0(t)\right) .$$

If the distribution is discrete long-concave (which is equivalent to (p_{k+1}/p_k) is non increasing) the counting function $\vec{\nu}$ is readily verified to be slowly varying. The fact that U is slowly varying (but does not satisfy the extended regular variation property) is well known [2].

For all $k \geq 1$, $p_{k+1} \leq \nu_1(0,p_k) \leq p_{k+1}/(1-p_{k+1}/p_k)$. If $p_j \geq \frac{1}{t} > p_{j+1}$, then $j \leq U(t) \leq j+k$ where $k \leq 2 + \log(p_j/(p_j - p_{j+1})/\log(p_j/p_{j+1})$ which is bounded. This entails that $\vec{\nu}(1/t) \leq U(t)$ and $\lim_{t\to\infty} U(t)/\vec{\nu}(1/t) = 1$. ■

In [10], the ETAC code escapes the $n+1^{\text{th}}$ symbol X_{n+1}, if $X_{n+1} \geq M_n$ where

$$M_n = \min\left(n, \{k \ : \ X_{k,n} \leq k\}\right) , \tag{13}$$

while $(X_{k,n})_{k\leq n}$ is the non-decreasing rearrangement of $X_1, \ldots, X_n$. The random threshold M_n is concentrated around m_n where $m_n = m(n)$ and $m(t)$ is defined for $t \geq 1$ as the solution of equation $U(t/x) = x$. In [10], it is proved that the function m inherits the regular variation properties of U. Namely, if $U(t) = \gamma^{\gamma} t^{\gamma} \widetilde{\ell}(t^{1+\gamma})$ where $\widetilde{\ell} = (1/\ell^{1+\gamma})^{\#}$ and $\ell \in \text{RV}_0$, m is regularly varying with index $\gamma/(\gamma+1) = \alpha$,

$$m(t) \underset{t\to\infty}{\sim} \gamma t^{\alpha} \ell(t) \underset{t\to\infty}{\sim} \frac{\alpha}{1-\alpha} t^{\alpha} \ell(t) \underset{t\to\infty}{\sim} \frac{\alpha}{1-\alpha} \vec{\nu}(1/t) .$$

When ℓ is slowly varying, the connexion between $\vec{\nu}(1/\cdot)$ and m is more subtle. The function m is the reciprocal of a De Bruijn conjugate of U. Hence by Lemma 24, in order to have $m(t) \underset{t\to\infty}{\sim} \ell(t)$, we need to have to have

$$\frac{\ell\left(\frac{t}{\ell(t)}\widetilde{\ell}_0\left(\frac{t}{\ell(t)}\right)\right)}{\ell(t)} \underset{t\to\infty}{\sim} 1 .$$

We are not aware of any meaningful characterization of this property.

APPENDIX D
NEGATIVE ASSOCIATION

When handling finite or infinite urn models, negative association arguments are a source of elegant moment or tail inequalities for occupancy counts. They complement Poissonization arguments.

Definition 25 (NEGATIVE ASSOCIATION)**.** Real-valued random variables $Z_1, \ldots, Z_K$ are said to be negatively associated if, for any two disjoint subsets A and B of $\{1, \ldots, K\}$, and any two real-valued functions $f : \mathbb{R}^{|A|} \mapsto \mathbb{R}$ and $g : \mathbb{R}^{|B|} \mapsto \mathbb{R}$ that are both either coordinate-wise non-increasing or coordinate-wise non-decreasing, we have:

$$\mathbb{E}\left[f(Z_A).g(Z_B)\right] \leq \mathbb{E}\left[f(Z_A)\right].\mathbb{E}\left[g(Z_B)\right] .$$

Theorem 26. *[18] For each $n \in \mathbb{N}$, the occupancy scores $(N_n^j)_{j\geq 1}$ are negatively associated.*

Monotonic functions of negatively associated variables are negatively associated. Hence, the variables $(\mathbb{I}_{\{N_n^j>0\}})_{j\geq 1}$ (respectively $(\mathbb{I}_{\{N_n^j=0\}})_{j\geq 1}$) are negatively associated as increasing (respectively decreasing) functions of $(N_n^j)_{j\geq 1}$.

APPENDIX E
CONCENTRATION AND MOMENT BOUNDS

Poisson distributions satisfy Bennett and thus Bernstein inequalities [See 9, Chapter 2].

Lemma 27. *For any Poisson distributed random variable N, for all $t > 0$,*

$$\mathbb{P}\left(N \geq \mathbb{E}N + t\right) \leq \exp\left(-\frac{t^2}{2(\mathbb{E}N + t/3)}\right),$$
$$\mathbb{P}\left(N \leq \mathbb{E}N - t\right) \leq \exp\left(-\frac{t^2}{2\mathbb{E}N}\right).$$

Bounds on expected inverse binomial random variables can be found in the literature [See 25, 3, and references therein]. The next results were developed for the purpose of this paper.

Lemma 28. *Let N be a binomial random variable with parameters n and p.*

$$\mathbb{E}\left[\frac{1}{N - \frac{1}{2}} \,\middle|\, N > 0\right] \leq \frac{1}{\mathbb{E}N}\left(1 + \frac{9}{\mathbb{E}N}\right).$$

Proof of Lemma 28: Using the fact that, for $k \geq 1$,

$$\frac{1}{k - \frac{1}{2}} = \frac{1}{k+1}\left(1 + \frac{3}{2k-1}\right) \leq \frac{1}{k+1} + \frac{9}{(k+1)(k+2)},$$

we have

$$\begin{aligned}
\mathbb{E}\left[\frac{1}{N - \frac{1}{2}} \,\middle|\, N > 0\right] &\leq \frac{1}{1-(1-p)^n} \sum_{k=1}^{n} \binom{n}{k} p^k (1-p)^{n-k} \left(\frac{1}{k+1} + \frac{9}{(k+1)(k+2)}\right) \\
&= \frac{1}{1-(1-p)^n}\left(\frac{\mathbb{P}\left(\mathcal{B}(n+1,p) \geq 2\right)}{(n+1)p} + \frac{9\mathbb{P}\left(\mathcal{B}(n+2,p) \geq 3\right)}{p^2(n+1)(n+2)}\right) \\
&\leq \frac{1}{np} + \frac{9}{(np)^2}.
\end{aligned}$$

■